CAPACITY ANALYSIS FOR MULTI-PRODUCT, PARALLEL-SITE ALUMINUM INGOT PRODUCTION

by

Brian C. Hintzman
M.S. Electrical Engineering, Stanford University (1991)
B.S. Electrical Engineering, Stanford University (1990)

Submitted to the MIT Sloan School of Management
and the
Department of Civil and Environmental Engineering
in Partial Fulfillment of the Requirements for the Degrees of

Master of Science in Management
and
Master of Science in Civil and Environmental Engineering

in conjunction with the

LEADERS FOR MANUFACTURING PROGRAM
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June, 1998

Signature of Author ______________________________ MIT Sloan School of Management
MIT Department of Civil and Environmental Engineering
May 8, 1998

Certified by ______________________________ Steve Graves, Professor of Management Science
Thesis Supervisor

Certified by ______________________________ Kevin Otto, Professor of Mechanical Engineering
Thesis Supervisor

Certified by ______________________________ Ismail Chabini, Professor of Civil and Environmental Engineering
Thesis Reader

Accepted by ______________________________ Joseph Sussman, Chairman of Departmental Committee on Graduate Studies
Department of Civil and Environmental Engineering

Accepted by ______________________________ Lawrence Abeln, Director of Masters Program
Sloan School of Management

CAPACITY ANALYSIS FOR MULTI-PRODUCT, PARALLEL-SITE ALUMINUM INGOT PRODUCTION

by

Brian C. Hintzman

Submitted to the MIT Sloan School of Management and
the Department of Civil and Environmental Engineering
on May 8, 1998 in Partial Fulfillment of the
Requirements for the Degrees of Master of Science in Management and
Master of Science in Civil and Environmental Engineering

ABSTRACT

This work focuses on casting aluminum ingots in a high volume production environment to supply an aluminum rolling mill. Ingot production is frequently the bottleneck activity and ingot availability has significant impact on all product lines. This research aims to increase the understanding of manufacturing throughput and opportunities to increase capacity. The scope of this analysis encompasses ingot production in two factories, separated geographically but producing to meet a single, combined demand. The products vary in two primary aspects: alloy composition and size (length x width x height). The research explores variation in demand for each product, rearranging the sequence in which products are manufactured, and adjusting which equipment is utilized to manufacture each product.

The approach uses a customized software tool developed specifically for this work through a third-party software vendor. The tool uses elements of both Discrete Event Simulation to capture critical sequence dependencies and Queuing Theory to provide production cycle times. This "hybrid" approach creates an equitable balance of these two methods to produce an accurate representation of production while keeping software execution times to an acceptable level (10-15 minutes per run). Results are validated by comparing output of the model with actual historical production of the ingot plant.

A Sensitivity Analysis is performed to determine the high leverage points in the manufacturing process. The inputs to the model are adjusted individually to determine their relative impact on overall production. In this way, all of the elements of manufacturing are ranked by their impact to production. This prioritization allows ingot plant engineers to verify these factors against intuition, select future capacity improvement projects based on their predicted impact, and evaluate potential changes in assignment of products to equipment.

Finally, a number of future production scenarios are discussed. This includes changes in overall demand, adjustments in product mix, and the results of capital improvements. The model is used to evaluate the predicted ingot capacity for each scenario. The modeling technique developed in this thesis provides improved results over other methods used in the past for similar activities.

Thesis Supervisors:	Steve Graves, Professor of Management Science
	Kevin Otto, Professor of Mechanical Engineering
Thesis Reader:	Ismail Chabini, Professor of Civil and Environmental Engineering

ACKNOWLEDGMENTS

The author owes deep gratitude to many individuals who made this work possible. I would like to thank, first and foremost, my supervisor at the company, Jeanette Nymon. She provided indispensable guidance and direction in my work. Perhaps more importantly, great patience during the struggle through many hurdles, discoveries, and setbacks necessary to make the model usable. I am indebted to the expertise of Matt Cesario, graduate of Stanford's SEMA program. His insight into, and understanding of, the manufacturing process was, and continues to be, critical to the application of this work. Also, Ken Schrock provided important background both on the details of ingot manufacturing as well as the sources of data within both ingot plants. He also provided unique perspective into the application of haiku.

There were many people in Factory A, where much of the research took place, who assisted with time and energy in this project. I would like to thank Dan Ritter, the ingot plant manager in Factory A for his support of this work. The ingot plant engineers, whose intimate knowledge of the ingot plant was crucial, include Paul Schaab, Jim Luebke, Dave Fenton, Jacqui Meenan and Julius Miller. Special thanks to the data keepers who were instrumental in pointing me in the right direction to find data: Rene Aller and Cindy Barnett. Paul Weber, the ultimate user for this work in the factory, for his assistance and letting me use his office for months on end. Rob Herman for his coaching and patience as well as Montie Gauss for reviewing the thesis. Finally, I would like to thank the management of the mill for their support and sponsorship of LFM work, Joe Velez and Ron Machetta.

There were several individuals in Factory B as well who were extremely helpful during the course of this work. These include most importantly Ken Williams who tolerated my several impositions upon visiting the site. Also, Uma Shankar who was very helpful in locating online databases and is one of the most experienced data warehousers I have encountered.

I would also like to thank the third party software vendor. They contributed not only the software tool and methodology but also insight into the types of data that are critical to building a capacity model. I would especially like to mention Greg Diehl and Christos Cassandras.

Finally I would like to thank my MIT advisors. They provided coaching and advice as well as patience with the progress of this work. Specifically, Prof. Steve Graves from the Sloan School of Management, Prof. Kevin Otto from the Department of Mechanical Engineering, and Prof. Ismail Chabini from the Department of Civil and Environmental Engineering.

Apologies to any and all individuals who also contributed I have inadvertently left out.

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

Chapter 1: Background

The focus of this work is the casting of aluminum ingots in a high volume production environment to supply several product lines in an aluminum rolling mill. Analyses in the past by the company sponsoring this research indicate that ingot production is frequently the bottleneck activity in the mill. Ingot availability has significant impact on the throughput of all product lines. The objectives of this work are to 1) Improve ingot capacity forecasting capability, 2) Examine future scenarios that provide opportunities for capacity improvement, and 3) Build the capability to examine two separate ingot plants as a single virtual factory model.

This chapter describes the environment in which the sponsoring company operates to provide a context in which to view the research that was performed. It also provides background describing the manufacturing operations for producing aluminum ingots. In concludes with the goals of the research which the rest of the thesis addresses in detail.

1.1 Industry

Aluminum is a metal characterized by its high strength-to-weight ratio. It is utilized in many applications for this reason. Common uses are for bike frames, beverage containers and, of course, the ubiquitous aluminum foil. It is also a primary material for aerospace applications including structural supports and the outer layer of aircraft. The latest applications are in the automotive industry where steel has long been dominant due to its lower cost. For comparison, recent prices for steel ingot are approximately $.07 per pound while aluminum ingot can cost $.70 per pound [5]. Several automotive manufacturers are exploring using aluminum to construct the frame of the car. Audi has already put this into production with its A8 automobile.

1.2 Factories

This work concentrates on two ingot plant facilities. Both facilities are part of the same business unit and thus work to satisfy the same product demand. However, the two plants are in different locations. This separation causes production coordination problems between the facilities. One factory, referred to as Factory A, combines an ingot plant and a rolling mill under one roof in the same facility. The second factory, referred to as Factory B, is a site which contains a smelting operation and an ingot plant in the same location.

1.3 Ingot Plant

The ingot plant produces aluminum ingots from pure aluminum and a variety of alloying agents. Ingots are blocks of solid metal containing a uniformly distributed mixture of aluminum and the alloying components (e.g silicon, iron, magnesium, etc.). The ingots studied in this work are rectangular blocks approximately 2 feet thick, 5 feet wide and 15 feet long. As pure aluminum has a density of 0.0965 lbs./ inch^3, a typical ingot weighs on the order of 25,000 pounds.

The two ingot plants considered in this study serve two fundamentally different purposes. For Factory A, the ingot plant is considered the beginning of the line in the manufacturing flow for the entire rolling mill. The main function of the ingot plant in this facility is to remelt scrap metal produced throughout the operation of the rolling mill and is thus referred to as a "cold metal" ingot plant. The rolling mill takes ingot from the two ingot plants to produce fundamentally two types of product: flat sheets, or plates, and coil. The next major processing step for the ingots after leaving the ingot plant is a hot rolling process, called the Hot Line. This step both flattens out and elongates the ingot, bringing it closer to the desired thickness and length for the final product. After rolling, the metal is either wrapped into a coil or cut into individual sheets in preparation for further processing. From there, the metal is sent to a job shop environment where additional processing steps are performed to achieve the final desired dimensions as well as thermal and mechanical treatments to obtain the metallurgical properties required by the customer. The entire rolling mill can be depicted in a simple, three-step block diagram:

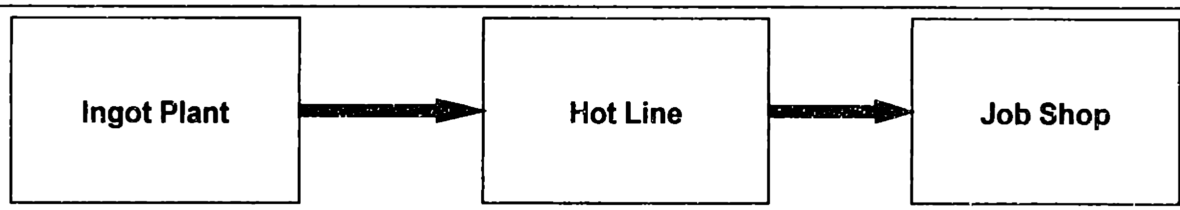

Figure 1.1 Simple Rolling Mill Block Diagram (Factory A)

Factory B, the ingot plant co-located with the smelting operation, casts molten metal produced by smelting. The entire factory is located in a region where electric power generation is cheap. This is important because smelting requires an enormous supply of electricity. Electric power is a primary cost component of smelting. In this facility, molten aluminum can be carried immediately from the smelting operation directly to the ingot plant in large, 12000 pound crucibles. This preserves the heat energy already imparted into the metal by smelting and eliminates the need to remelt the metal in the ingot plant. Since the ingot plant receives all of its metal for casting already in molten form, it is referred to as a "hot metal" ingot plant. A simple block diagram of the facility is shown here:

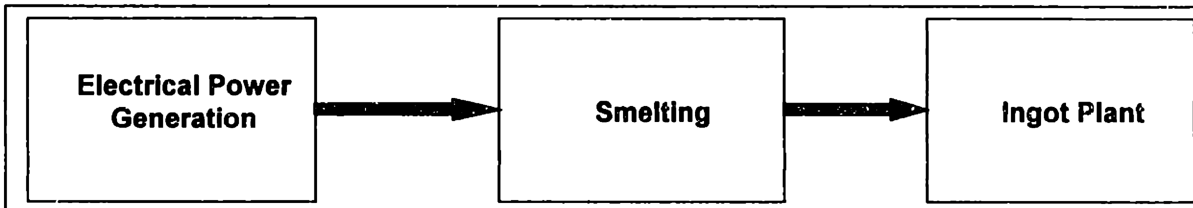

Figure 1.2 Simple Smelter Block Diagram (Factory B)

1.3.1 Products

The ingots produced by the plants are differentiated by two key factors: alloy and dimension. The alloys, as described above, are determined by the relative concentration of aluminum and a number of alloying metals. The dimension of the ingot is determined by the size of the cross-sectional mold and the length to which it is cast.

1.3.1.1 Alloys

The aluminum industry has established a classification system for designating alloys by composition. This allows a common reference both for manufacturers and customers. The classification system is a four digit system where the first digit denotes the major alloying element.

Designation	Primary Alloying Element
1***	None
2***	Cu
3***	Mn
4***	Si
5***	Mg
6***	Mg and Si
7***	Zn

Table 1.1 Aluminum Alloy Designation System [4]

The alloy composition is extremely important to the customer. Producing larger number of alloys impacts capacity because a significant amount of time is required when a casting changes from one alloy to another. During the time of this study there were 70 unique alloys in use in the ingot plants. A table is used by the ingot plant engineers that determines the severity of the alloy transition for each of the 4900 possible combinations of old alloy and new alloy.

1.3.1.2 Dimensions

The other major product differentiation is dimension. As a three-dimensional rectangular block, this consists of gauge, width, and length. However, manufacturing considers only cross-section and length because gauge and width are not determined independently. An ingot mold is a two-dimensional form that establishes the gauge and width together for the product (see Figure 1.3). The length is determined independently from the mold. The cross-section dimensions are important because of the capacity impact when cross-sections change. This is similar to the effect of alloy changes as discussed in the previous section.

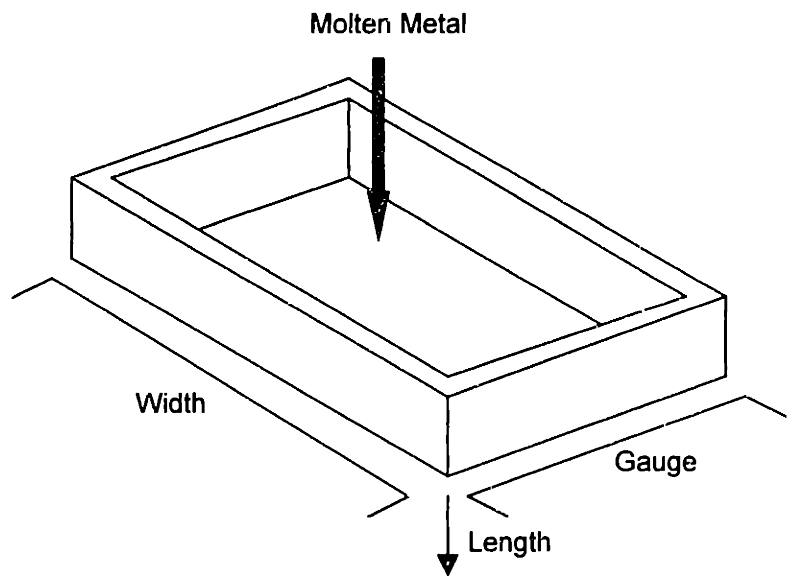

Figure 1.3 Ingot Mold

1.3.2 Equipment

The equipment used in the ingot plants is physically very large. Typical ingot batch sizes are on the order of 100,000 pounds and standard aluminum molten temperatures are between 1200°F and 1300°F. The equipment therefore must be built to handle this weight and heat 24 hours a day, 7 days a week, essentially indefinitely. The capital expenditures required for even one casting Complex cost many millions of dollars. In addition, because of the enormous physical size of the equipment, modifications to and movement of the machinery happens very rarely. Due to these factors the configuration of an ingot plant changes very slowly and installed equipment remains largely unchanged for decades.

The equipment of an ingot plant is grouped into elements called "Complexes". A Complex, at the simplest level, is made up of a furnace and a pit. The furnace stores the molten alloy metal until 1) a sufficient amount of the metal is available to cast the desired number and size of ingots, and

2) the casting Pit has been cleared and is ready for the next cast. Block diagrams for the equipment in both factories are presented in Figures 1.4 and 1.5 below.

1.3.2.1 Furnaces

Cold metal ingot plants, Factory A in this study, require two furnaces for each casting Complex. The first furnace is referred to as the Melter and is labeled (n)M in Figure 1.4 where (n) represents the number of the Melter furnace. As its name suggests, this furnace is used to melt all scrap material necessary for one cast. This collection of metal is called the Charge, and consists of pure aluminum and the constituent alloying elements. Hot metal ingot plants do not require Melter furnaces. Factory B derives its raw material directly from smelting and the metal arrives already in molten form. The metal is simply weighed as it arrives and based on the alloyed contents of the incoming metal it is directed to furnaces that require additional material.

There is a second type of furnace used in the plants referred to as a Holder. The Holder is labeled (n)H in both Figure 1.4 and Figure 1.5 where (n) represents the number of the Holder furnace. The Holder receives its metal from the Melter, in the case of cold metal plants, via a trough that pours metal directly from one furnace to the other. In the hot metal plants, the metal arrives at the Holder from the smelter in molten form. The Holder furnace stores molten metal while it is being prepared for the cast and provides heat energy to keep the material in the molten state. Metal is then poured directly from the Holder during the cast utilizing both Holder and Pit.

Equipment in the ingot plant changes slowly due to the high capital costs. Investment in Factory A over the last 40 years has resulted in two fundamentally different kinds of furnaces. Three new Complexes were installed that substantially increased output (the group labeled West in Figure 1.4). All operating parameters in these furnaces are electronically monitored and controlled from control houses, one per Complex. Most significant is that the Holder furnaces actually tilt on hydraulic cylinders during the cast. Since this action is monitored by computer it provides much greater operational control over the flow rate of the metal or to stop the flow if necessary. It also means that nearly all of the metal can be poured out of the furnace during the cast. The sizes of the objects in Figure 1.4 indicate the relative furnace size for each Complex as well as the relative number of ingots that can be cast by each Complex.

Another group of Melters and Holders has been in place for several decades (the group labeled East in Figure 1.4). The equipment is substantially older and more manual. This is also the style of equipment used in Factory B, shown in Figure 1.5. More importantly, the furnaces are much smaller in capacity than the newer furnaces which limits the overall capacity. In addition, the Holder furnaces do not tilt as the new equipment does. Metal is poured from the furnace during a cast simply by removing a metal plug near the bottom of the furnace and allowing gravity to drain the metal. Not only must this be manually replaced in case the cast must be aborted, but this leaves what is called a "heel" of metal in the furnace below the height of the plug. This remaining metal becomes scrap during an alloy change and represents a much larger percentage of the original Charge than is the case with the new furnaces.

It is important to note the furnace configurations associated with 5PIT and 6PIT in Factory A and 1PIT and 2PIT in Factory B. These Complexes have two Holders adjacent to the single Pit. This was necessary due to the equipment capability at the time of installation. Because of the thermal output and overall size of furnaces available at that time, it was determined that the casting Pit would normally be idle awaiting the Holder to be prepared for the next cast. Therefore, in order to improve asset utilization on the Pit, two Holders were installed for each Pit in these Complexes. This allows the two Holders to alternate casts from one side of the Pit to the other utilizing much of the same equipment. This configuration is referred to as a Dual-Furnace Pit.

While this configuration is sensible from a capital expenditure viewpoint, it makes equipment utilization measurement and capacity analysis much more difficult.

1.3.2.2 Pit

The casting Pit is the final equipment stage for forming the ingot. As the molten metal comes from the Holder to the Pit, it is directed into a separate drain for each ingot. Through this drain the metal enters the two-dimensional mold depicted in Figure 1.3. This form contains the metal only as long as it is molten. The metal is supported by a hydraulically controlled platen that carries the weight of the ingot as it is cast. As the metal enters the mold, the metal that touches the mold freezes, forming a hard shell around the outside of the molten material. When the cast starts the platen begins to drop at a predetermined rate for the alloy, eventually moving the cast material past the bottom of the mold.

At this point the frozen aluminum shell continues to act as a mold for the molten metal in the center of the ingot. As this process continues the ingot begins to grow in the vertical direction. The length of the ingot is then determined by the amount of time the supporting platen spends moving downward. A corresponding amount of molten metal is poured during this time so that the ingot reaches its desired size. Because the ingot is physically moving downward during the cast, the procedure is commonly referred to as a Drop.

Once the cast is completed and the ingots are of the prescribed length, the molds are removed from what is now the very top of the ingots. The ingots remain in the Pit to cool for an amount of time that varies by alloy. An overhead crane then grabs each ingot individually and takes it to be weighed. This is known as stripping the pit. After all of the ingots have been removed, the Pit is readied for the next cast.

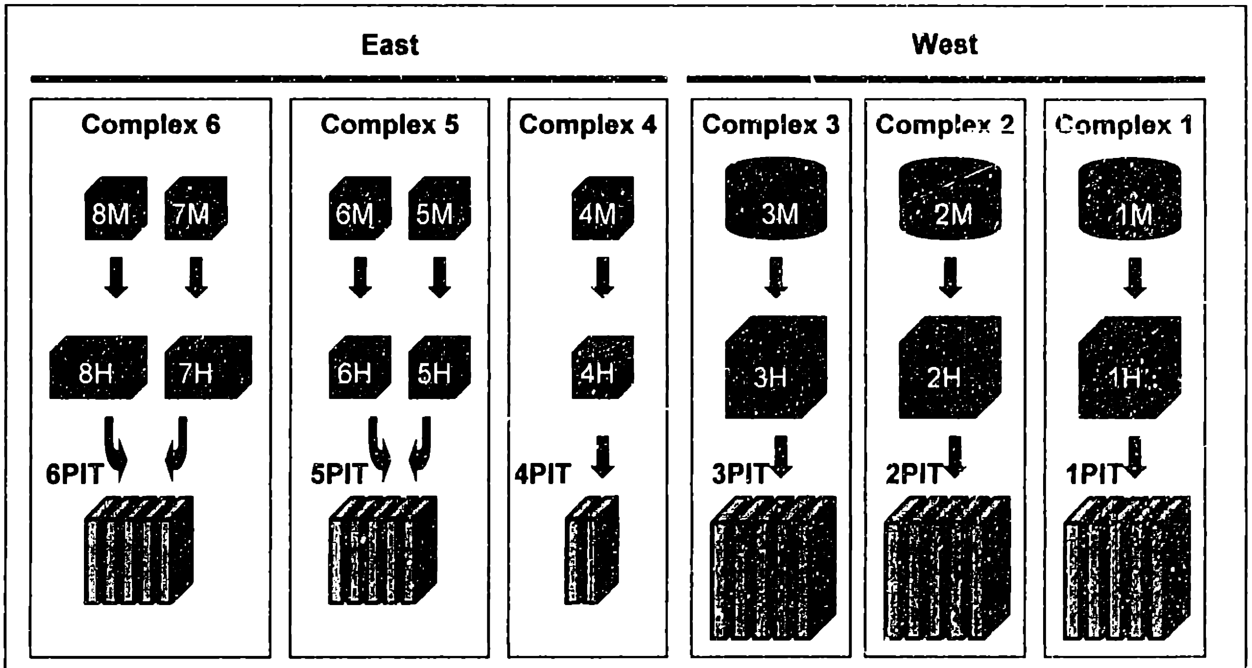

Figure 1.4 Ingot Plant Configuration for Factory A

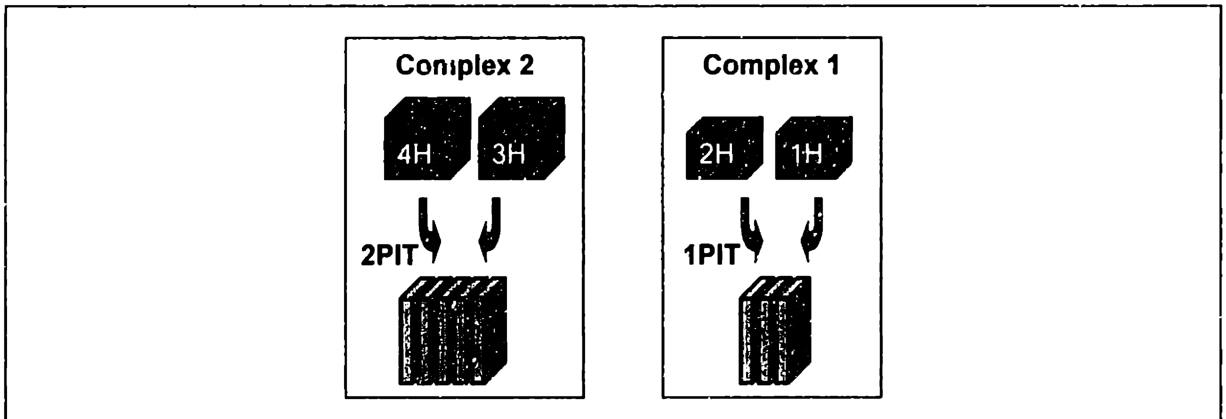

Figure 1.5 Ingot Plant Configuration for Factory B

1.3.2.3 Tooling

The term "tooling" is used in the ingot plant to refer to a collection of molds. The ingot mold tooling is assembled in the mold shop by mounting several molds onto a mold mounting plate. The number of molds attached varies by Pit from two to six. This determines the number of ingots that can be cast simultaneously. Pits with smaller Holders cast fewer ingots (2-3) because of the limited amount of total metal, whereas Pits with larger Holders can cast as many as six ingots at a time. All molds in a single tooling piece always have the same cross-sectional dimensions. Because the platen supporting the ingots is also a single piece, each ingot is cast to the same length for the Drop. Therefore, ingots cast in the same Drop always have the same overall dimensions.

There are two different styles of ingot mold tooling in use in the two factories: standard tooling and quick change tooling. The standard (older) style tooling takes two to three hours to change the Pit from one size to another. Newer quick change tooling is largely automated and can usually be changed in approximately 45 minutes.

Because of the setup time required to change tooling, Dual-Furnace Pits normally cast by alternating between each Holder without changing the tooling between each cast. On the occasions when the tooling is changed, the next cast from each Holder uses the new tooling size. Therefore, the two Holders can be treated as synchronized from the tooling standpoint. Both Holders change tooling from one size to another simultaneously.

1.4 Description of Problem

The subsequent customer of both ingot plants is the rolling mill associated with Factory A. It is managed with a flow path methodology that aggregates demand by major product category where six have been defined for this mill. Management of all flow paths begins with ingot supply. In recent years, as demand for products from the mill has increased significantly, it has typically been the case that the bottleneck of production lies in the ingot plants. Improved capacity forecasting in the ingot plants is needed to alleviate this problem. The company wants to understand the ability of the ingot plants to meet customer demand as production changes.

1.4.1 Current Practice

The current capacity forecast process begins in the marketing organization. The total demand for the coming year is compiled by marketing representatives for each product line. This is translated into a total demand for ingots in terms of alloy and dimensions. The ingot plant utilizes one or both of two available methods to interpret this demand. One method is spreadsheet based that encompasses historical production rates, average product yield, and expected planned and unplanned downtime. This produces an overall forecast for total number of pounds of expected production over the next year. This is used as a baseline for production commitments to the customer. The other method is a physical system where each cast of ingots is represented by a card. Many cards are arranged representing the order in which ingots will be cast with estimates for the time consumed by each cast. Estimates are also made for the time taken for changes of alloy or tooling.

The company would like to improve the capacity forecasting process. These models are deficient because they do not take into account variation in production cycle times or downtime. In addition, the existing practices do not provide the ability to model the tradeoffs in alloy transition time as the types of alloys vary. In general, both above approaches are static models that could be improved by dynamic modeling techniques.

1.4.2 Goals

There are three fundamental objectives to this work performed during a research internship with the sponsoring company. The first is to improve the annual capacity planning process for the ingot plant. The present process described above is believed to be, at best, accurate to within +/- 10% in predicting future capacity when compared to actual results. The ingot plant engineers would like to be able to analyze and assess new mixes of product and be able to accurately predict the plant's ability to meet customer demand. A tool that is able to model operational variation and improve on the accuracy of the currently available methodology would meet this goal.

The second goal is the ability to perform scenario analysis. Achieving this goal allows the ingot plant management to examine situations that have not occurred in the past. This includes new product assignments, such as manufacturing product types on equipment where they have never previously been produced. This would also allow management to understand how new equipment additions or upgrades to existing equipment would effect overall production. This capability would complement the understanding and intuition of ingot plant engineers of their operations into new, untested realms.

Finally, a longer term goal is to create a "Virtual Ingot Plant" model. Such a tool would combine the capacity forecasts for both Factory A and Factory B in a single model. This capability would allow the mill operations to better understand the total supply of ingots that can be utilized throughout the rolling mill. For the ingot plant management, it adds the ability to consider both switching product types between the two factories as well as which ingot types to outsource outside of the factory network. This objective was not completed during the research. The final chapter of this thesis discusses why it was not finished, what is required to complete this work, and expands on the benefits that a virtual factory model can provide.

1.5 Chapter Summary

This chapter introduced both the industry environment as well as the manufacturing context for this research. It also described the goals and scope of the research as agreed upon with the sponsoring company. The remainder of the thesis describes how the research proceeded in order to fulfill these goals.

Chapter 2: Modeling Approach

Over the past ten to fifteen years the company sponsoring this research has built several ingot capacity models. These models of ingot plant operations were built both by internal engineers and, in some cases, in conjunction with outside consultants. Several of these methods are still in use to some degree as discussed in the previous chapter. Other models have been abandoned and replaced by new, superior models. From an accuracy and usability standpoint, the new model used in this research was developed with two primary goals in mind: it must account for process variability and it must execute in acceptable length of time.

The approach chosen by the sponsor company for this effort was a partnership with an independent software vendor with specific expertise in manufacturing throughput modeling. This partnership produced two distinct pieces of software as the result of this work. In this thesis, the words "software", "tool", and "software tool" are used synonymously to refer to the application program built by the independent software vendor. This tool and the algorithm it utilizes are the subject of Chapter 2. The words "model", "database", and "database model" are used synonymously to refer to the collection of data that characterizes production in the two ingot plants in this study. This model was built by the thesis author with the assistance of many individuals at the internship company. Its construction is the subject of Chapter 3.

These two pieces of software, the software tool and the database model, cannot operate without each other. When both pieces are referred to together in this thesis they are jointly called the "modeling tool". To further clarify this issue, an analogy can be drawn with the widely used application Microsoft Excel. Microsoft Excel is analogous to the software application built by the independent software vendor. A speadsheet used in Microsoft Excel is analogous to the database model built by this author. For confidentiality reasons, the internship company asked that I not use the name of the software tool in this thesis.

The tool uses elements of Discrete Event Simulation to capture critical sequence dependencies created by the order in which products are made, decisions which are made in production scheduling and on the factory floor. The software also utilizes Queuing Theory to estimate production cycle times, including both processing time and time waiting for equipment to come available. This "hybrid" approach is used to provide an equitable balance between these two methods in order to accurately model production while keeping software execution times to a satisfactory level (10-15 minutes). This algorithmic approach is called Aggregate Dynamic Modeling (ADM) [8]. The details of this approach are described in this chapter.

2.1 Interface

The software is designed to run on a Windows-based PC. It requires that Microsoft Access be available on the PC. The Microsoft database package is used to manipulate the potentially large amounts of data necessary to characterize and describe manufacturing processes and store the resulting output calculations produced as a result of running the model. The input and output data (described in the following two sections) can be viewed and manipulated separately in Microsoft Access; however the capacity simulation can only be performed in the software tool developed for this purpose.

2.1.1 Inputs

The software requires a complete description of the manufacturing process in order to estimate the number of ingot pounds which can be produced. All major input categories are listed and described in Table 2.1. These inputs can be entered either directly through a user interface of the software or by opening the database in Microsoft Access and entering the information there.

Input Type	Description
Forecast Period	The total number of days available in ingot production analysis period.
Equipment	The model has a number of predefined Complex definitions, including all those available in Factory A and B (i.e. with and without Melter, one or two Holders). For each piece of equipment the user specifies the expected scheduled and unscheduled downtime used to model equipment availability.
Labor	Definitions are supplied for the labor groups available in the ingot plants, the Complexes the labor groups are assigned to operate, and the percentage of time the labor groups are unavailable and limit production (e.g. meetings, absenteeism, etc.).
Tooling	The number and types of mold cross-sections that are available for each Pit are entered. This includes the number of ingots that can be cast simultaneously and the percentage of time that the tooling is unavailable due to repairs.
Process Steps	For each alloy cast at each Complex, all processing steps are listed that required to produce ingots. This includes the sequence of the steps, the average and coefficient of variation of cycle time, the percentage of material that is scrapped, and equipment and labor that is utilized.
Alloy Transition Process	This defines the processing required to convert a furnace from an old alloy to a new (described in detail in Chapter 3). Each process includes cycle time and pounds of scrap produced.
Alloy Transition Matrix	This matrix lists the required alloy transition process required for every combination of old and new alloy in the ingot plants.
Demand Per Product	This provides the total number of pounds required of each ingot type during the ingot production analysis period.
Product Assignment	This information specifies which Complex(es) will be used to manufacture the product. An entry is required for each product-Complex pairing as well as a percentage of product pounds that are to be manufactured there.
Casting Plan	There is a Casting Plan for each furnace specifying the order in which products are to be manufactured. This includes the sequence of alloys and cross-sections that are to be manufactured.

Table 2.1 Inputs required to construct ingot capacity model

2.1.2 Outputs

The outputs of the model are generated by the software after running the simulation. These outputs indicate many aspects of the way in which the ingot plant is utilized.

Output Type	Description
Complex Information	The model generates a significant amount of information for each Complex over the forecast period including total pounds produced, total pounds melted, average yield, average cycle time, etc. The most important statistic of this information utilized in this research is Percent Complete (discussed below).
Asset Utilization	For each major piece of equipment (Melter, Holder, Pit) the model produces an asset utilization calculation representing the percentage of total time in the forecast period when the equipment is actually working on good product. The remaining equipment time is broken down into categories including equipment downtime, setup time, blocked/starved time, and yield loss time (time spent on bad product).
Product Information	This information is very similar to that presented in the Complex Information, however, it is broken down by product rather than by Complex. It includes forecast pounds of good production, total pounds melted, yield, etc.
Tooling	For each piece of tooling, the model calculates the percentage of time it is actually utilized and the average number of Drops that are delayed because the tooling is not available (in use on another pit or in repair).
Labor	Labor time is calculated and categorized by active time, idle time, and absenteeism time.

Table 2.2 Outputs generated by ingot capacity simulation

In this study, the primary outputs of interest are the Percent Complete calculation (presented as part of the Complex Information) for each Pit and the idle time (presented as part of the Asset Utilization information) for each piece of equipment. For clarity, the formula and interpretation of Percent Complete is provided in Figure 2.1:

$$\text{Percent Complete} = \frac{GoodIngotPoundsMade}{TotalIngotPoundsDemanded} * 100$$

Percent Complete < 100% ➔ Complex over loaded, cannot complete orders in forecast period
Percent Complete ≅ 100% ➔ Complex fully or under loaded, finishes all orders.

Figure 2.1 Equation and Interpretation for Percent Complete

Using the Percent Complete calculation, quick inspection determines the Complexes that the model calculates to be over loaded, those with Percent Complete below 100%. For Complexes where the Percent Complete number is equal to or just above 100%, the user must dig further to examine the Asset Utilization calculations. For these Complexes, if the equipment also shows significant idle time, the model has determined that not only did the Complex complete all of the ingots demanded, but it also had significant additional time where the equipment was simply idle with no demand. The algorithm does not artificially create additional demand in under loaded situations so it is impossible for the Percent Complete to rise much above 100%. Close examination of both of these numbers is necessary for an initial understanding of model results.

2.2 Software Tool Algorithm Description

The algorithm used by this software is a hybrid of both Discrete Event Simulation and Queuing Theory techniques. These two methods are used within the tool to produce a single set of results to the user.

Discrete Event Simulation (DES) of a manufacturing system simulates every event that occurs in the manufacturing process. This includes each process step on each piece of equipment in the proper sequence as prescribed by the process. Cycle times are assigned from probability distributions that characterize each process. The advantage of this method is that it can be very detailed by simulating every event that occurs in manufacturing instead of using approximations. However, the primary disadvantage of this method is that it is extremely compute intensive.

The Queuing Theory portion of the model is added to retain maximum accuracy while significantly reducing the required execution time. The primary savings results from removing the need to generate a discrete probabilistic cycle time for every event in the model. Cycle times are created as an aggregation of Queuing Theory results comprising both queuing time waiting for equipment to become available as well as the actual processing time on the equipment. An overview of the calculations are provided in the Queuing Theory discussion below.

2.2.1 Discrete Event Simulation of Ingot Plant Operations

The DES portion of the tool centers around the Casting Plan, the user specified sequence of alloys and cross-sections for each furnace. The engineers in both Factory A and B have developed over time a heuristic method for assigning alloys to run on each furnace. This is based primarily on grouping alloys in order to minimize setup time penalties when transitioning from one alloy to another in the furnace. The Casting Plan must also be sensitive to physical restrictions in the ingot plant. There is a small number of alloys that, due to differences in furnace equipment, must run on certain furnaces. In addition, because not all ingot molds sizes have been purchased for each Complex, not all ingot cross-sections are available on all Pits. The user of the software must be aware of these limitations in the ingot plant when creating the Casting Plan.

The Casting Plan is detailed at the "Campaign" level. A Campaign is defined as any number of sequential Drops that are performed without changing either the alloy or the cross section used. The user therefore lists the sequence of all Campaigns expected to occur during the forecast period. In addition, the model has the ability to generalize a repeating manufacturing cycle in order to simulate a long forecast period. In such cases, the user specifies a sequence for a subset time period (e.g. 8 weeks for a 52 week forecast period). The tool then simply replicates this sequence a sufficient number of times to fill the forecast period.

The software uses the Casting Plan to create a detailed sequence of every Drop needed to fulfill the demand for each product in the forecast period. The Demand Per Product input includes information provided by the user for average Drop size in pounds and Campaign size in pounds. The algorithm uses this to determine how many total Drops overall are required to produce the specified demand for each product as well as the typical number of Drops that are performed in sequence before switching to another alloy or cross-section. Finally, note that because the Casting Plan only includes alloy and cross-section, each entry in the Casting Plan can correspond to multiple products because a length is not specified. So without changing alloy or cross-section, demand for several products of different length can be satisfied. An example of the detailed Casting Plan process is depicted in Figure 2.2:

User-Specified Casting Plan
Each entry is a Campaign

Alloy	X-Section
A	G1-W1
A	G2-W2
B	G2-W2
C	G1-W1
C	G3-W3

Model-Generated Detailed Casting Plan
Each entry is a Drop

Alloy-Gauge-Width-Length
B-G2-W2-L1
B-G2-W2-L1
B-G2-W2-L2
B-G2-W2-L3
B-G2-W2-L3

Figure 2.2 Example of assignment of products to Campaigns

In building the detailed Casting Plan the algorithms checks to see if each Campaign is required. There either may be no demand for the products corresponding to the Campaign or Campaigns earlier in the Casting Plan may satisfy all the necessary demand. Skipped Campaigns generally mean one alloy or one tooling transition will not be necessary. The detailed Casting Plan therefore determines the complete sequence of events during the forecasting period including tooling changes, alloy transitions, and the severity of alloy transitions due to the sequence of alloys. Each entry in the detailed Casting Plan represents one Drop.

Finally, the case of dual furnace pits adds more Complexity. Each furnace has its own sequence of alloys because the contents of the furnaces can be processed independently. However, the two furnaces must be coordinated in terms of tooling since they cast into the same pit. This means that, for the simulation and in reality, not only do both furnaces lose availability of the pit during a tooling change but also both furnaces change tooling at the same time. This means that the cross-sectional dimensions of the products coming from both furnaces must be synchronized.

2.2.2 Queuing Theory Applied to Ingot Plant Operations

Once the sequence of products is determined, the algorithm can begin assigning cycle times to these events. Cycle times are simply the processing times for each step necessary to produce an ingot as recorded by equipment operators. These values are then incorporated both with the variability of cycle times, provided by the user, and the queuing effects created by utilization of the equipment At a high level, the mathematics of the queuing theory portion of the modeling algorithm rests on five basic equations shown in Figure 2.3:

1. $T = S + R * L$

2. $X = \frac{D}{L} * \frac{1}{A}$

3. $U = X * T$

4. $W_{AVG} = F * X$

5. $F = \left(\frac{C_a^2 + C_s^2}{2} \right) (W_{ARR.} * T) + T$

Where:

Variable	Description	Dimensions
S	= Setup Time per Lot	[Minutes/Lot]
R	= Run Time per Drop	[Minutes/Ingot]
L	= Lot Size in Ingots	[Ingots/Lot]
T	= Equipment Time per Lot	[Minutes/Lot]
D	= Demand in Ingots	[Ingots/Forecast Period]
A	= Avg. Percent Equipment Availability	[Dimensionless]
X	= Number of Lots Processed	[Lots/Forecast Period]
U	= Utilized Equipment Time	[Minutes/Forecast Period]
F	= Flow Time per Lot (Queued + Process Time)	[Minutes/Lot]
W_{AVG}	= WIP in Lots in System (Equipment + Queue)	[Lots]
$W_{ARR.}$	= WIP in Lots in Queue (Lot arrivals)	[Lots]

Figure 2.3 Fundamental Queuing Theory Equations for the Tool

The outcome of these equations is a value for the flow time, or the time a lot spends waiting and being processed by a single piece of equipment. This result is shown in Figure 2.4:

$$F = \left(\frac{C_a^2 + C_s^2}{2}\right)\left(\frac{U}{1-U}\right) * (S + R * L) + (S + R * L)$$

Where:

$W_{ARR.} = \dfrac{U}{1-U}$ Approximation for G/G/1 Queue, random arrival of orders

U = Utilized Equipment Time

C_a = Coefficient of variation for arrival processes

C_s = Coefficient of variation for service processes

Figure 2.4 Queuing Theory Result Equation for the Tool

The equations shown above are solved through an iteration process. The key to this process is determining the values to be plugged in for C_a and C_s. These values are solved utilizing a proprietary set of differential equations developed by the software vendor. However, this is the step that includes all known process variation, including cycle time, labor availability, and tooling availability (described in Sections 2.2.3.1-2.2.3.2). The resulting value for flow time can, in turn, be used for time consumed both waiting and processing in each of the three pieces of equipment in the Discrete Event Simulation. The equation for arrival of WIP assumes that when orders arrive, they find the system in a random state, rather than a known state where either the equipment is idle waiting for the next order or is always busy processing the previous order.

The advantage of this approach is that the flow time values do not need to be computed for every ingot Drop simulated in the forecast period. Notice that because the variables shown in Figure 2.3 are values for the overall forecast period (e.g Demand in Ingots, Number of Lots Processed) the calculations are aggregated at the forecast period level rather than at the individual event level. These values can be calculated once for each unique ingot type (alloy and cross-section).

The values can then be inserted into the DES. In this way, the queuing theory result is aggregated over the entire forecast period. Individual queue sizes and queuing times are not calculated for each event in the discrete event simulation. This provides enormous reductions in the amount of computation required and therefore the total execution time experienced by the user.

2.2.3 Aggregation of Discrete Event Simulation and Queuing Theory

Finally, the action that binds the DES and Queuing segments together is an iteration process. The DES portion executes creating the order of Drops, the Queuing section assigns cycle times to the determined order of operations, and additional adjustments are made for labor and tooling availability as described below. This process then loops for 20 to 30 executions, internally recording output results for each iteration. Because the detailed Casting Plan is filled with a pseudo-random assignment of Drops, the results are different on each execution. These results are then averaged to produce the final output results. This allows the model to account for process variation and aggregation of the ultimate results.

2.2.3.1 Labor

The assignment of labor to production operations occurs in the model after the DES and Queuing segments have run. Labor groups are assigned using a "Processor Sharing" algorithm. In time-sliced computer mainframe operating systems, executing jobs are prioritized based on the inverse of cycle time. Using this method, jobs with short execution time will tend to experience short queuing time while waiting to execute. This mirrors the way in which tasks are prioritized by workers in the ingot plant. In cases where a short amount of time in needed to keep the casting process running, individuals will stop what they are doing, take care of the short task needing attention, and return to the previous job. So, again, short jobs experience short waiting time. The Queuing mathematics treat this type of system as a G/G/n/k queue, where n is the number of workers in the labor group and k is the number of machines they cover.

The assignment of labor groups allows the software to determine both the amount of cycle time consumed in waiting for labor as well as the percent of labor utilization reported in the output results. The total manufacturing throughput is then reduced for the Complex based on the labor waiting time required. The final results then take account for the assignment of labor.

2.2.3.2 Tooling

The utilization of mold tooling is handled in a similar fashion. After accounting for labor assignments, the software calculates the amount of tooling utilization based on the number of Drops that use each tool size. But this calculation is used to determine asset utilization only and is not used to adjust total output. This is done because the experience of the ingot plant engineers is that operations are almost never constrained by tool availability. In cases where either a tool is being used in another Pit or is down for repairs, orders are reordered to avoid the problem and keep all Pits running. The software simply assumes that excessive tooling utilization will be handled by intelligent product scheduling. This is not assumed for labor because personal absences cannot be anticipated and often do impact factory operations.

2.2.4 Unique Software Issues

There are several other issues fundamental to the operation of the ingot plant that necessitated the development of a custom software approach. The first is an issue that has been touched upon already: the sequence dependent effect on throughput.

2.2.4.1 Alloy Transitions

This is most evident in the handling of alloy transitions. The amount of time necessary to prepare a furnace to process a new alloy is directly dependent on which alloy was produced previously. Depending on the elemental components of the two alloys in sequence the amount of setup time necessary varies from none at all to a shift or more of usable production. A software tool is needed that not only understands this issue but is capable of storing the huge number of possible alloy combinations. At a lower level, the tooling transitions present a similar issue. The software must fundamentally understand the tooling associated with each product. With this knowledge it can determine whether production time needs to be consumed performing a tooling change from one product to another.

2.2.4.2 Equipment Coupling

The second issue is for the software to understand the close coupling of the Holder and the Pit. In many manufacturing flows, products flow from one machine to the next for discrete processing steps. In this case, the metal receives unique processing in the Holder alone (alloying), both the Holder and the Pit (casting), and the Pit alone (stripping the Pit). It therefore must be possible in the software to assign equipment independently to each batch as it flows through the Complex. It is then possible to properly simulate the coupling of the Holder and Pit during the time that this equipment is consumed on the same batch of product, and allow the equipment to separately process separate batches at other times.

2.2.4.3 Finite Buffering

The third topic is closely associated with the previous discussion. From a queuing theory view of the equipment, the design of the Complex is a system with finite buffering. Specifically this means that there is no queue in front of each piece of equipment, the queue is the previous piece of equipment and that queue can only have a size of 0 or 1. When processing is completed at one stage, the downstream equipment is checked. If it is available the batch moves along, if it is not available the batch must remain where it is. This not only continues to occupy the equipment it also prevents any upstream batch from moving forward. This is necessary because the batches of metal are so large and require so much heat to remain in the molten state that it is not economically viable to construct additional furnaces simply to act as buffers. The processing equipment must therefore also be used as buffers for the Complex. Standard queuing packages have trouble with this issue and are typically limited in the types of finite buffering they handle.

Any one of these issues alone could be handled with standard Queuing mathematics. However, it is the combination of all three of these issues together that necessitates the approach that has been taken. The modeling challenge can be viewed as a Markov Chain problem where the software would compute the probability of the system transitioning from one state to another. It was determined that the state space would blow up with a huge number of possible states. The DES portion makes this problem manageable by pre-determining the sequence of operations. This allows the software tool to know apriori: 1) Where are we in the Casting Plan?, and , 2) Is each piece of equipment running, coupled, starved, or blocked?

Since a custom package was going to be required to handle the mathematics, the software was designed with several other benefits. For instance, the package implicitly understands the different equipment configurations (with and without Melter, single vs. dual furnace Pits). There is a facility for specifying Casting Plans along with the alloy and cross section information. The software also understands the alloy transition concept and provides both a means to define the process and an assignment of processes to alloy pairs.

All of the issues described led to the conclusion that commercially available capacity modeling tools were not sufficient to handle the specific needs of ingot production. There is not only sufficient need within the sponsoring company for such a tool but also sufficient specific ingot plant requirements to justify the software package. This, therefore, is the tool applied in this research.

2.3 Applications

Finally, it is worthwhile to discuss those areas in which the modeling tool described in the previous sections can be appropriately applied. The tool is built with predictive capabilities in mind as well as the ability to examine production modifications that may affect overall capacity. However, the choice of modeling approach also imposes certain restrictions and specifically does not address several important issues that are presented below.

2.3.1 Capabilities

The aggregation inherent in the tool limits its application for short time frames where specific downtime issues or outlying cycle times are more likely to affect the accuracy of model prediction. The tool is most appropriate in long time frame scenarios, generally three months or more, where long run averages increasingly reflect typical, historical production output.

2.3.1.1 Capacity Forecasting

The first, best application of this tool is for forecasting future capacity situations. The user of this tool may enter future demand as anticipated on a per product basis. In addition, these products must be assigned to casting Complexes where they are expected to be produced. Given this information, the software predicts the total number of pounds of each ingot type that can reasonably be expected to be produced based on past performance. This output can be expected to be more accurate than current, static forecasting methods.

2.3.1.2 Opportunity Evaluation

A second and useful application of the tool is for understanding the major manufacturing leverage points in the factory. This is developed by performing a sensitivity analysis on the key input components that comprise the capacity model. As the tool produces overall Percent Complete values, each input component can be individually modified to determine the extent to which that variable impacts total output. This allows ingot plant engineers both to reinforce intuition as to the importance of each factor but also rank order areas of future improvement in terms of the anticipated percentage gains in output. The ingot plant obtains maximum value from future capital expenditures to increase capacity when prioritized using a quantitative approach. Sensitivity analysis using the model is presented in Chapter 4.

2.3.1.3 What-If Analysis

Finally, the development of this methodology is designed to create an ability to perform "What-If" analysis. This allows the user to simulate situations that do not presently exist in the plant. These include primarily equipment additions and upgrades as well as new assignments of products to casting Complexes. Whereas sensitivity analysis prioritizes those factors with the largest direct impact on manufacturing output, What-If analysis goes a step further. This method produces actual estimates for the total amount of production that can be expected in new scenarios.

This capability allows the ingot plant to simulate new modes of plant operation without having to spend valuable equipment time and effort actually acquiring a statistically significant amount of production data. It also gives engineers the opportunity to examine a larger number of product assignments than is feasible in the actual operation. With a rapid modeling capability, the intent is to provide the capability to examine the maximum number of situations in the minimum amount of time on the part of ingot plant engineers. Several What-If scenarios are described in Chapter 5.

2.3.2 Limitations

As in any application, there are limitations to the range of situations in which the tool can reasonably be expected to be applied. The primary limitation is in time scale. As an Aggregate Dynamic Model, the software tool is best applied in long time frames where long run average behavior adequately describes behavior. For time frames as short as a day, for instance, only one to five Drops may actually be completed. The tool is typically inaccurate in such time scales as this is not a sufficient number of Drops over which to average typical variability in cycle times. The software is best used in time frames no smaller than one month. This "aggregated" behavior impacts several more issues described below.

2.3.2.1 Dynamic Behavior

It is normal for manufacturing organizations to be in a continual state of flux. Continuous improvement, or Kaizen, has become a major focus of American corporations in the 1990's. For manufacturing, this results in the reduction of cycle times and the elimination of process steps. This is a complication to modeling efforts because the process becomes a moving target. Characterization of the process requires that cycle times are static during the time period of interest.

The usefulness of this tool is limited to time frames during which the manufacturing process is largely unchanging. Additions of equipment or the removal of major processing steps during the forecast period of the software cannot be accommodated. Major events effecting manufacturing have to be treated as two piece-wise contiguous time periods: the time before the change occurred and the time after. Each time period must be simulated separately with two models and the results combined by the user to produce data for the desired time period.

2.3.2.2 Scheduling

This tool is not capable of creating a detailed, daily production schedule for the ingot plants. It cannot examine the demand for each ingot type and produce a production sequence that maximizes overall output. This would require ordering the sequence of ingot dimensions to minimize the total number of required tooling changes during the time period of examination.

This tooling constraint would also have to be balanced with the need to sequence alloys to minimize the time required to perform alloy changeovers. An optimization procedure would need to produce a separate sequence for each casting Pit. Such an algorithm was not requested by the sponsoring company and was not built by the software vendor.

Another factor limiting the applicability of the tool to scheduling is its inability to take into account the needs of timely delivery of product. The algorithm described above would produce a sequence of products with very large batch sizes such that each product type is only manufactured once during the forecast period in sufficient quantity to satisfy all demand. This would minimize the number of alloy and tooling transitions necessary. However, while this type of scheduling may maximize production, it may not be optimal for the rolling mill as a whole.

If, for instance, the forecast period is one year, this type of scheduling would result in large amounts of inventory as the products would be manufactured in many cases long before they were needed in the rolling mill. In the past few years, the overall rolling mill operation has greatly increased its focus on inventory reduction and delivery time performance. The ingot plants have responded by building to order the products needed most immediately by the rolling mill to satisfy customer demand. This has resulted in smaller batch sizes and more alloy and tooling transitions as ingot plants produce to fulfill immediate demand rather than store large amounts of finished ingot. This prioritization of manufacturing goals decreases rather than increases total capacity.

2.3.2.3 Optimization

The final issue to note is that in several significant ways this type of software is not an optimization tool. Most importantly, the algorithm does not optimize production across Pits. This means that if the way in which the user has assigned products to Pits results in a Pit is that is, for example, significantly under utilized, the tool does not automatically shift products from other Pits to consume the unused capacity. Each Pit is examined separately, calculating the total production for each Pit based on the products assigned to that Pit. The balance of products across Complexes is dependent upon the percentage of output assigned to each Pit and therefore, ultimately, the user.

Another optimization opportunity would be to prioritize selected alloys based on their profitability. In cases where ingot demand exceeds capacity, a sensible approach would be to produce as much volume as possible of the most profitable alloys and outsource the rest of the demand. Ideally, the user would specify the profit margins for each alloy thereby imposing a priority for the alloys. The software would report the quantity of the high margin alloy demand that could be completed as well as the amount of lower margin alloys that must be outsourced. Such a facility has not been requested and is not presently built into the tool.

2.4 Chapter Summary

This chapter presented the software tool built by a third party software vendor for this research. It described, to the extent possible without violating confidentiality concerns, the algorithm implemented in the tool to determine ingot plant capacity. It further presented some advantages and limitations of the software approach that was chosen for this project. This software package provided the means to model and analyze the manufacturing operations of ingot production that is described in the remainder of this thesis.

Chapter 3: Building an Ingot Plant Model

Given the background software approach, this research utilizes this software tool to model the ingot production environment, evaluate the model, and make recommendations. The first step was to construct the ingot plant models. This consumed the majority of the research effort in order to locate the appropriate data for the many information needs of the software and to ensure that metrics used in the factories are expressed in the form expected by the tool.

This chapter describes the process of building ingot plant database models to use in conjunction with the software tool presented in Chapter 2. It describes how data for the major input components of the model were obtained. It concludes by describing the method used to validate the models against actual production at the two ingot plants in this study.

3.1 Model Construction

Building the database model required gathering the appropriate data to develop all of the inputs described in Table 2.1. The following sections describe the data gathering methods and types of data utilized to address each of the major input components. As the methods applied in Factory A and B largely overlap in approach, Factory A is used for discussion and examples except in those areas where the two factories differ significantly.

The primary purpose for this software is for use as a capacity forecasting tool. However, before forecasting could be addressed, a model had to be built and validated for prediction accuracy as measured by Percent Complete (see section 2.1.2). The initial model construction efforts were focussed solely on this validation task. The approach was to build a model representative of a past historical period and verify that the production estimates produced by the software matched what was actually produced during that time period.

3.1.1 Equipment

The only equipment pieces defined in the model are the major elements comprising the Complex: Melter, Holder, and Pit. The Complex types are chosen first (e.g. MHP, MMHHP, etc) in the software tool. This instructs the software how many pieces of equipment to expect in the database model and of which type. The primary information required for the equipment is scheduled and unscheduled downtime in order to determine availability.

Unscheduled downtime is entered in the format of Mean-Time-To-Fail/Mean-Time-To-Repair (MTTF/MTTR). This was calculated from approximately one year worth of maintenance records for the equipment. This is both a sufficient time period for consideration and the equipment outages during this time are typical of future downtime expectations. Scheduled downtime events were not included in the MTTF/MTTR calculations. For the validation work, actual downtime periods were accumulated for the historical period. For use as a forecast tool, anticipated (scheduled) equipment maintenance and upgrades must be considered. Since all equipment is required to manufacture an ingot, scheduled downtime for the entire Complex was accumulated and applied entirely to the Melter (Holder for Factory B), implying that the entire Complex is unavailable.

3.1.2 Labor

The work crews in the ingot plants are described in the model as labor groups. The model information includes the number of workers in each crew, the average percentage of time they are unavailable, and the process steps during which labor is required. The number of workers assigned to each crew and the steps for which they are needed is readily available. However, there was little data as to the percentage of time that labor availability actually affects production. This was due not only to record keeping but also to a flexible and dynamic workforce where worker substitution can occur freely. The overall labor unavailability was assumed to be 5%.

In addition, the overhead cranes that load the Melter with cold metal and a separate set used to remove ingots from the Pit are also modeled as labor groups. This is done because the cranes can be utilized by several Complexes, similar to work crews, and the cranes are only required for certain process steps, as is the case for standard labor. Maintenance records were again used to estimate availability.

3.1.3 Tooling

A list is maintained by the ingot plant engineers of ingot mold tooling sizes available at each Complex along with the number of slots available for ingots. However, the software requires this information as well as ingot mold tool unavailability data. Ingot mold downtime is supplied in two parts: 1) Number of Drops between maintenance (MTTF), 2) Average time to repair (MTTR). There is no record in the ingot plants of the amount of time when production is limited by ingot mold tooling repairs. It is the consensus of the ingot plant engineers that tooling availability is almost never a constraint to production. Therefore the average time to repair is set to 0, thereby ignoring this factor.

The other factor describing tooling is the time necessary to change the Pit from one tooling type to another. The primary factory here is whether a Pit has been upgraded to quick change tooling. This allows the Pit to change to a new tool in about 1/3 of the time required by older equipment generations. The tool change time in practice is therefore independent of tooling size. Average times were then entered for each Pit depending on whether it utilizes quick change tooling.

3.1.4 Process Steps

The definition of process steps is much more involved than any of the factors discussed so far. The software tracks a unique set of process steps for each combination of alloy and Complex.

This allows the necessary flexibility due to significant processing differences for varying alloys as well as differences in cycle times between old and new equipment.

3.1.4.1 Process Flow

The level of detail at which the process flow is disaggregated and measured is markedly different at the two factories. The most significant difference is the fact that Factory B does not require Melter furnaces and therefore does not need the associated process steps. Further, the process steps for which cycle times are recorded vary in detail. These are listed in Figure 3.1 along with check marks indicating the equipment that is utilized during each step. In addition, the steps are listed so that those most functionally similar in the two plants are horizontally aligned:

Factory A

Operation	M	H	P
Charge	✔		
Melt	✔		
Skim	✔		
Alloy/Stir	✔		
Sample1	✔		
Transfer	✔	✔	
Stir		✔	
Sample2		✔	
Flux		✔	
Sample3		✔	
Salt/Skim		✔	
Preset		✔	✔
Cast		✔	✔
Term			✔
Pull			✔
Setup			✔

Factory B

Operation	H	P
Fill	✔	
Process	✔	
Cast	✔	✔
Turnaround		✔

Figure 3.1 Process Steps in Factories A and B

The function of the major process steps are described here:

- Fill: During this time the furnace is filled with metal. In Factory B, crucibles of molten metal are carried from the smelter and poured directly into the Holder furnace. In Factory A, the charge of scrap metal is loaded into the Melter furnace, the metal is completely melted and any contaminate particles are skimmed off the surface, the composition is sampled to confirm correct alloying, and the metal is transferred to the Holder.

- Process: While the metal is in the Holder, additional alloying material may be added, a flux step may occur where argon gas is bubbled through the melt to trap additional contaminants, the liquid is skimmed again and prepared for casting.
- Cast: Clearly this is the time when the metal is actually being poured from the Holder into the Pit. This time will vary based on length of the ingots and the Drop rate of the alloy, as described previously.
- Turnaround: Finally the ingots are pulled from the Pit and the Pit is prepared for the next cast.

In addition to the sequence of process steps and the assignment of equipment to each step, the cycle time distributions for each process step is specified. The cycle time distributions are described by the mean cycle time and coefficient of variation. This information was obtained from cycle time sheets recording the number of minutes consumed by each operation over a three month period. This data was used to calculate the average and coefficient of variation for the cycle times. Due to confidentiality concerns the actual cycle time estimates are not presented in this thesis.

3.1.4.2 Branching

Another element of process definition is the ability to define branching operations. These are operations that are not performed on every batch of a given alloy. These include operations that are up to the judgment of the operator or may be a function of cleanliness of the furnace, therefore purity of the alloy. The software tool allows the user to specify a percentage of time this operation is performed. When there is no branching, this number is simply 100%.

This issue was difficult to characterize in the production environment. As the branching operations are often up to the discretion of the Complex operators, there may be no discernable pattern of branching. Therefore, the percentage of branching used in the model was based on the judgement of the ingot plant engineers. We used discretion in deciding whether to include branching effects in the process definition. In general, branches performed more than 10% of the time (either taken or not taken) were coded as branches rather than as always or never performed.

3.1.4.3 Scrap

The final component of process definition is describing the amount of scrap generated at each step. This enables the tool to determine the amount of metal that each batch must begin with in order to complete with the necessary number of pounds for each Drop. Based on the record keeping in the ingot plant, scrap comes from two sources: melt loss and scrap loss. Melt loss is generated during the melt process. It is comprised of metal that actually burns or evaporates away as well as metal that clings to the floor and sidewalls of the furnace after the metal moves on to the next piece of equipment. Scrap loss can be thought of as the metal that is actually cast and then scrapped. This consists of ingots that crack or have some physical defect, ingots that do not meet the elemental requirements for the alloy, and casts that begin and then are aborted due to equipment failure.

Values are calculated for melt loss by comparing the weight of metal that goes into the initial furnace (Melter in Factory A, Holder in Factory B) to the weight that is cast. The difference in weight is simply the aggregate metal that is lost between the start of the batch and the cast. This is calculated as a percentage of the total initial weight of the batch. The scrap loss is measured directly as scrapped ingots are weighed and categorized as scrap. This is tabulated in a computer database. For this model, the melt and scrap loss values were calculated separately for each alloy.

In the case of Factory A, the values were further separated between the old and new equipment, or East and West sections of the factory (see Figure 1.3).

3.1.5 Alloy Transition Process

An alloy transition describes the process necessary to use a furnace to cast a new alloy different from the alloy most recently cast. This step is critical to understand and characterize because, depending on the severity of the transition, it can be the longest single operation that occurs in the ingot plant. The alloy that was last cast from a furnace is called the "From" alloy and the new alloy to be used is called the "To" alloy.

The severity of transition is determined by the largest difference in percent concentration of any single alloying element when comparing the composition of the From alloy to the To alloy. There is never a concentration problem for elements where the concentration increases when moving from the old alloy to the new alloy. A sufficient level of the alloying agent is simply added to the new charge in order to meet the specifications. When the concentration must be reduced, however, a calculation must be performed to determine whether enough pure aluminum can be added to the furnace (without overfilling) to sufficiently dilute the element to meet the concentration limits of the new alloy.

The database model describes alloy transitions in terms of the time necessary to perform the procedure and the number of pounds of scrap metal generated during the process. A series of interviews were performed with ingot plant personnel to determine the best way to organize and describe the many functions that are performed during transitions. In actual production, there is essentially a continuum of subtle differences in operations depending on the two alloys involved, the amount of metal remaining in the furnace at the end of the last cast, the amount of time since the last maintenance on the furnace, and many other factors. To simplify the procedures sufficiently so that they could be entered into the model, it was agreed that Factory A uses three primary classes of alloy transitions. These transition classes were assigned the names Conversion, Drain&Clean, and Drain/Clean/Wash.

Table 3.1 describes the average time and scrap pounds associated with each transition class. In order to clearly describe this table the counting variables *i*, *j*, and *k* are used. *k* represents the number of Pits in the plant (six for Factory A). *j* represents the number of transition classes used in the plant (three for Factory A, described in the previous paragraph). *i* represents the number of operations that can be performed during an alloy transition (number of rows in Table 3.1).

In Table 3.1, the Equipment column lists the equipment types that must be serviced and the Operation column lists all operations that can be performed during a transition (there are *i* operations). Next there is a column for each transition class (there are *j* transition classes). The numbers in these columns represent the fraction of time, F_{ijk}, that each operation is performed for each transition class on each Pit. The remaining columns show the time, T_{ijk}, to perform each operation and the pounds of scrap, S_{ijk}, produced for each operation, grouping similar Pits where appropriate. The actual time and scrap values are not provided for confidentiality reasons.

Equipment	Operation	Conversion	Drain&Clean	Drain/Clean/Wash	1-3 Pit Process Time (Minutes)	4-6 Pit Process Time (Minutes)	1-3 Pit Scrap (Pounds)	4 Pit Scrap (Pounds)	5 Pit Scrap (Pounds)	6 Pit Scrap (Pounds)
Filter Box	Alloy Box	1	1	1						
	No Recycle Flush	0	0.25	0	T1	T2	S1	S2	S3	S4
	Recycle Flush	0.95	0.75	0.9	T3	T4				
	Wash Cast	0	0	1						
Melter	Alloy Furnace	1	1	1						
	Drain/Clean	0	1	1	T5	T6				
	Wash	0	0	1	T7	T8				
Holder	Alloy Furnace	1	1	1						
	Drain/Clean	0	1	1	T9	T10	S5	S6	S7	S8
	Wash	0	0	1	T11	T12	S9	S10	S11	S12

Table 3.1 Alloy Transition Time and Scrap for Factory A

Using Table 3.1, a simple calculation derives the cycle time and scrap produced for each transition class for each Pit. To calculate the total alloy transition time for a Conversion, for example, the average time for every operation is summed together. The average time for each operation is found by multiplying the fraction, F_{ijk}, by the actual cycle time for that operation, T_{ijk}. The total amount of scrap produced is calculated in the same fashion, substituting scrap in place of cycle time for each operation. These calculations are summarized in Figure 3.2

For (Pit)$_k$, (Transition Class)$_j$:

$$\left.\begin{aligned} (\text{Total Cycle Time})_{jk} &= \sum_i F_{ijk} * T_{ijk} \\ (\text{Total Scrap Pounds})_{jk} &= \sum_i F_{ijk} * S_{ijk} \end{aligned}\right\} \text{Where } i \text{ specifies unique equipment operations}$$

Figure 3.2 Calculation of Alloy Transition Time and Scrap Pounds

A similar interview process was performed with ingot plant personnel in Factory B where a more complicated situation exists. Due to differences in available equipment, as well as developed practices in the plant, it was agreed that four distinct alloy transition classes are in use in this facility: Conversion, Crucible Flush, Filter Box Vacuum, and Drain/Clean/Wash. The last three classes each have two variants, listed in Table 3.2, depending on the alloys involved. The result of these interviews produced a characterization table for Factory B depicting cycle time and scrap produced. It is presented in Table 3.2.

Finally, these values are used to characterize the alloy transition classes using the same method as for Factory A (see Figure 3.2). However, for the three classes with two variant procedures, the two variants are weighted equally (each term in the summation is multiplied by ½) in the calculation of total cycle time and scrap produced for the transition class.

Equipment	Operation	Conversion	Cruce Flush w/o scrap	Cruce Flush w/scrap	Filter Box Vacuum "Easy"	Filter Box Vacuum "Hard"	Drain&Clean	Drain/Clean/Wash	Process Time (Minutes)	1 Pit Scrap (Pounds)	2 Pit Scrap (Pounds)
Filter Box	Alloy Box	1	0	0	0	0	0	0			
	Cruce Flush w/recycle	0	0.3	0	0	0.6	0.3	0.7	T1		
	Cruce Flush w/o recycle	0	0	0.7	0	0	0	0	T2	S1	S2
	Vacuum	0	0	0	0.4	0.6	0.3	0.7	T3	S3	S4
	Wash Cast	0	0	0	0	0.6	0.3	0.7	T4	S5	S6
Holder	Alloy Furnace	1	0.3	0.7	0.4	0.6	0	0			
	Drain/Clean	0	0	0	0	0	0.3	0.7	T5	S7	S8
	Wash	0	0	0	0	0	0	0.7	T6	S9	S10

Table 3.2 Alloy Transition Time and Scrap for Factory B

3.1.6 Alloy Transition Matrix

The final element that is independent of forecast period is the alloy transition matrix. It is simply a list of every combination of From and To alloy along with the transition class (described in the previous section) associated with that alloy pair. As noted in Chapter 1, this includes 4900 alloy pairs. This information was also entered into the model. It is used by the DES portion of the calculation to determine the amount of cycle time consumed when performing alloy transitions.

3.2 Model Validation

With all of the time-period-independent information collected, actual validation of the model required selecting an actual historical time frame to compare. A time period of three months was selected. The final pieces of data were then assembled as described below.

3.2.1 Forecast Period

The only information required here is the number of days included in the forecast period. As the months chosen were March though May, the total number of days is 92. This informs the tool how much total time is available for production in order to meet the demand.

3.2.2 Demand Per Product

For an historical time frame, the demand is actually the total number of pounds of good ingots for each product produced during the time frame in question. This is convenient because this is a number that is largely known. Two difficulties were encountered during this research. First, one database that records pounds produced by the ingot plant also records other downstream operations but does not clearly differentiate the operation. Therefore, total pound data becomes corrupted with additional pounds weighed in at other points in the process flow. A second, separate database records production pounds along with ingots that are returned to the ingot plant long after they are cast. This occurs when ingots are found to have defects during downstream operations. These returns reduce the total number of pounds of production. Unfortunately, by the time the ingots leave the ingot plant, it is impossible to know which Complex produced the ingot. Understanding that no available data was perfect, several corrections were made to the best available data to get a number close to the actual number of pounds for each product.

3.2.3 Product Assignment

Building on the demand data, the model also needs to know the assignment of percentage of pounds to each furnace. The tool uses this to determine how the product mix is actually distributed to each furnace. For this, it is necessary to go a step further than simply knowing the demand for each product. It is necessary to know the pounds of each product produced on each furnace. Again, using data corrected to the extent possible (as described above) we calculated percentage numbers for this category to add to the model.

Another component of the product assignment data specifies average Drop size in pounds and average campaign size in pounds. Using this information, the software tool knows how many pounds of good production to expect from each Drop for each product. This defines the fundamental batch size in pounds, unique to each product. For reasons discussed in Chapter 4

(Sensitivity Analysis), the average campaign size was set equal to the average Drop size. The Sensitivity Analysis determined that the average campaign size information does not impact the Percent Complete calculation.

3.2.4 Casting Plan

The final piece of the database model is entering the Casting Plan. As described in Chapter 2, this is the expected sequence of alloys and cross-sections manufactured in each furnace. This information is readily available in a validation context because the exact sequence of production is known. This was simply obtained from ingot plant records and entered into the model.

3.3 Validation Results

As stated at the beginning of the chapter, most of the research work was consumed in building the model and getting the results to accurately reflect historical plant output. Many of the initial models showed output that was significantly different from what had been achieved in the past. A long investigation period was necessary to calibrate the model, relying heavily on sensitivity analysis work described in the next chapter. The final validation results were as shown in Figure 3.3:

Factory A	1PIT	2PIT	3PIT	4PIT	5PIT	6PIT
% Complete	100.8%	91.5%	95.7%	97.3%	95.0%	96.7%
Factory B	**1PIT**	**2PIT**				
% Complete	101.3%	96.2%				

Figure 3.3 Validation Results for Factory A and B

Recall from Chapter 2 that the Percent Complete calculations the result of primary interest. It measures the percent of ingot pounds demanded of each Complex that the software tool estimates can be produced during the simulation time frame. These results show that the modeling effort is reasonably accurate. In Factory A, for example, 1PIT is extremely accurate with the predicted production coming to within 1% of actual production. 2PIT, however, shows the worst result as the model predicts that only 91.5% of actual production can be achieved. The model estimates that the capacity for this Pit is lower than what has actually been produced.

Forecast periods of both three months and five months were examined to compare accuracy over differing time frames. Results varied somewhat over these time frames but generally stayed within the +/- 10% target of accuracy.

3.4 Chapter Summary

This chapter presented the methodology used to build the two ingot plant models that were the focus of this research effort. The primary input components of the model were described along with the process for obtaining the data to characterize those components. Finally, the validation methods were presented describing how the software tool was used with the database models to achieve capacity forecast results that matched actual plant capacity.

Chapter 4: Sensitivity Analysis

Sensitivity analysis can assist a model builder in understanding the modeling results. Sensitivity analysis indicates which of the input variables are most important in effecting the output results. Using sensitivity analysis, we can determine the high leverage points in the manufacturing process in order to prioritize the highest impact projects for increasing capacity.

This chapter describes the approach used to perform sensitivity analysis using the software tool and database models described in Chapters 2 and 3. First, the numerical method used to obtain sensitivity values is presented. This is followed by a description of the most important learning that was obtained by analyzing the sensitivity results.

4.1 Numerical Approach

Many modeling methodologies have clear, well understood procedures for performing sensitivity analysis. The Simplex Method for solving Linear Programs, for example, produces shadow prices that indicate the relative importance of each of the input variables. These values come "for free" as part of the optimization algorithm used to solve the program. In a similar fashion, the multiple regression technique for mapping a number of independent input variables to a dependent output variable produces several results including coefficients for the independent variables. These coefficients can be used to rank order the explanatory variables to determine those that have the most influence in explaining the behavior of the outcome.

For this modeling approach used in this research, however, there is no simple scheme that automatically produces sensitivity results. Not only is the algorithm a custom approach, as described in Chapter 2, but also the software is not designed to compute its own sensitivity analysis and produce values as part of its internal calculations. Therefore, performing a sensitivity analysis requires a "brute force", numerical approach.

As discussed in Chapter 2, the primary output of interest in the tool is the Percent Complete calculation. The sensitivity analysis determines the effect of each input variable on this output. Figure 4.1 describes the equation that is the basis of this method. The *%Complete* output can be considered as a function of the vector $\bar{x}$, where $\bar{x}$ is comprised of all the individual input variables, x_i. The effect on *%Complete* is measured by making changes to the input variables, Δx_i, and noting the change in %Complete. As Figure 4.1 shows, the vector $\bar{y}$ represents the base case of *%Complete* about which the changes to the input variables are made.

$$\%Complete(\bar{x}) = f(\bar{x})$$

$$\%Complete(\bar{x}) = \%Complete\big|_{\bar{y}} + \sum_i \frac{\partial f}{\partial x_i}\Delta x_i + \text{(Higher Order Terms)}$$

Figure 4.1 Computation of *%Complete* as a function of all input variables

Note that higher order interaction effects of two or more input variables were not examined and, hence, ignored. Because the software tool is an aggregate model as opposed to a precise discrete event simulation the higher order interaction estimates would not only be inaccurate but of little added value. The purpose here is to determine which easily definable and separable inputs have the most effect on production.

Next, a Sensitivity Coefficient is defined in the context of this research. It is most easily thought of as the percent change in output over the percent change in input. Each Sensitivity Coefficient is calculated by performing a separate run of the modeling tool with only a single input changed from the base case, $\bar{y}$. Therefore the overall effect on output, *Δ%Complete*, can be approximated by the linear combination of all of the changes in inputs multiplied by their respective Sensitivity Coefficients. This is shown in Figure 4.2.

$$\frac{\partial f}{\partial x_i} \cong (SC)_i$$

Where:

$(SC)_i$ is the Sensitivity Coefficient for each input x_i

Therefore:

$$\Delta\ \%Complete(\bar{x}) \approx \sum_i (SC)_i \cdot \Delta x_i$$

Figure 4.2 Effect of single variable sensitivity on %Complete

Because rerunning the model for every input is time consuming and considerable effort on the part of the user, the further simplification is made that the effect on *%Complete* is linear. Therefore a single value for the Sensitivity Coefficient is sufficient rather than attempting to construct a complete response surface for the effect of each individual input. These single values for the Sensitivity Coefficients are presented in Tables 4.1 and 4.2 at the end of this chapter.

This approach to producing sensitivity results presents one difficulty in that the *%Complete* value is discontinuous. Beyond *%Complete* values of 100%, the *%Complete* number stops rising and equipment idle time begins rising. When there is sufficient capacity to complete product demand during the forecast period, the tool attempts to "guardband" the results, meaning it produces slightly more output than necessary to protect against variability in production. When *%Complete* actually exceeds 100% the user of the software can be reasonably confident that it will be possible to complete the demand in actual production. During the course of this research we observed that the algorithm in the tool tends to peak at *%Complete* values of ~108%. Beyond this buffered output, the idle time in the asset utilization calculation begins to increase. This represents excess capacity available during the forecast period to produce more output.

This behavior in the model has two impacts on sensitivity analysis. First, the analysis must be performed in regions where *%Complete* is below 100%. This avoids the region above 100% where the correlation between input and output falls off rapidly. In some cases this means modifying the model to put *%Complete* in this region. The simplest way to accomplish this is to reduce the forecast period so that it is impossible to complete demand in the time available. Second, the user must take care to modify input values that tend to decrease total capacity. This ensures that not only is *%Complete* below 100% but also it will move farther below 100% as the input value changes.

4.2 Sensitivity Results

Using the approach described above, sensitivity analysis was performed on both Factory A and Factory B. The results for each Pit in the two factories are included in Tables 4.1 and 4.2 included at the end of this chapter. The Sensitivity Coefficients that are most important in effecting the output have been set in bold type. The results and discoveries from this analysis are discussed below.

4.2.1 Factory A

The sensitivity analysis for Factory A was performed during the course of this research primarily because of the difficulty involved in validating the initial model. The analysis proved to be a very useful tool in determining the areas in which to focus which has the greatest impact on *%Complete*. Perhaps an even bigger benefit was derived from determining those areas in which *not* to concentrate. Because of the data integrity issues in a number of areas, the analysis was able to eliminate the need to clean up input data for variables that had little to no impact on overall output

4.2.1.1 Significant Inputs

Based on the results of the analysis, a number of conclusions were drawn about those factors that are and are not important for the total capacity of Factory A. While there are many input variables, those that proved to be most significant, and in some cases surprising, are discussed below. In some cases, these also affected how the factory model was constructed.

- Melter Cycle Time: The fact that Melter Cycle Time was one of the most significant inputs for every Pit in Factory A was a confirmation both of the modeling algorithm as well as the numerical approach to sensitivity analysis. Based on past history as well as the knowledge of the plant engineers, the Complexes are typically Melter constrained. The time needed to melt down the charge of cold metal to a molten state dominates the other major activities in the Complex. It is standard to finish a Drop, which empties the Holder, and then wait a few minutes to an hour for the Melter to complete. Therefore changes in Melter Cycle Time have a very direct impact on overall throughput in a Complex.

- Drop Size: The Drop Size addresses the number of pounds that are cast in each Drop as function of ingot dimensions. This seems an obvious result as the total number of pounds produced is directly affected by number of pounds in each batch, as one would expect. In addition, at first inspection this seems a trivial calculation, given the dimensions of the ingot and the density of aluminum (.0965 $lbs/inch^3$, on average, varies slightly by alloy) one can calculate the approximate number of pounds in each ingot.

  However, the number of ingots cast in each Drop can vary. This is determined primarily by the number of positions available in the tooling and the number of pounds of metal available in the Holder. Typically, the operation will cast as many ingots as there are tooling positions. However, for large cross-sections or long ingots, there may not be enough metal available in the Holder to cast the full complement of ingots. In these cases the number of ingots per Drop are reduced. It was found that in the initial model, the average Drop size in pounds was seriously flawed. First, due to corrupted input data, the calculated average values were significantly incorrect. Then, a set of general rules were developed, replacing the calculations, relating Drop size to dimensions for every Pit.

What is important here is that this factor was identified by the sensitivity analysis as an input with highest significance. After the sensitivity analysis was performed these rules were refined with much finer precision because of continued problems with model validation. While there were potential problems in many areas of the model, the rank ordering of input variables provided by sensitivity analysis enabled further modifications of the model. This analysis proved to play a decisive role completing the model validation.

4.2.1.2 Non-Significant Inputs

Equally important in the course of model development was the segregation of input variables that have little to no impact on output. There were a number of areas where the data was ambiguous (data that was either unavailable, corrupted, or incomplete). The sensitivity analysis allowed the developer to determine those variables where guesses or approximations were "good enough". This eliminated the need to spend additional effort collecting hard-to-find data where results within an order of magnitude were sufficient. While again, a number of these issues were identified, two of the more interesting are discussed here.

- Campaign Size: This input determines the number of successive Drops that are made of any one ingot type when it is cast. Initially, considerable effort was expended in calculating this value in conjunction with Drop size, where many of the same difficulties were encountered.

After performing the sensitivity analysis, however, it quickly became apparent that this factor bore little importance.

Additional discussions with the software developer revealed why this should be the case. In the Discrete Event Simulation discussion in Chapter 2, the user-provided casting plan was explained. Recall that the software creates the detailed casting plan defining every Drop needed to fulfill demand (also see Figure 2.2). A further detail of this process is that tool attempts to fill *every* user-defined Drop before creating additional Drops. This most closely resembles actual production methods in the plant. Otherwise, the simplest algorithm is to fill all of the demand for each product the first time it appears in the user-supplied casting plan, which is specifically *not* how the ingot plant operates (also see discussion in Section 2.3.2.2)

Using this method, the tool itself determines the number of successive Drops (Campaign size) necessary to fill the required demand. This means that it is not necessary for the user to calculate average campaign size from historical data. Based on this insight, the campaign size field was changed in the model to simply one Drop, resulting in negligible change in *%Complete*.

- Double Pours: Another initially ambiguous issue was the handling of Double Pours. This activity occurs for small size ingots when the quantity of metal consumed by the Drop is one-half or less of the total metal in the Holder. In these cases, it is possible to cast a second time without requiring additional metal from the Melter. As the Melter is typically the bottleneck equipment, this procedure can usually be done with one Melter cycle. So, for example, where a single Drop might be of five ingots, the Complex operators may have the option to pour two Drops of three ingots each. This produces six ingots instead of five with no additional impact to capacity.

Significant effort was expending on this issue, not only determining the proper method for handling this issue but also in collecting accurate data. However, the sensitivity analysis showed that this issue, too, did not significantly effect results. When the processes defining these Double Pour events were removed from the model, *%Complete* barely changed. This is due to the fact that only some ingot dimensions have the potential for Double Pours, when they occur they add only one or two ingots, and they can occur on a maximum of half of the Drops (every other pour). All of these factors limit the applicability of the Double Pour procedure. Based on the sensitivity results, the decision was made to exclude these events from the model, resulting in a considerable simplification.

4.2.2 Factory B

The sensitivity results for Factory B are, not surprisingly, similar to those for Factory A. Specifically, the two Pits of Factory B perform comparably to Pits 5 and 6 of Factory A, because each of these pits have two Holders from which to cast into the Pit. Of course, Factory B does not have Melter furnaces so the number of input factors to explore is reduced.

4.2.2.1 Significant Inputs

First, it is important to note that, again, the Drop Size in ingots is the single most important factor in determining output, a completely reasonable result. The fact that it is significant in the same magnitude as is the case for Pits 5 and 6 in Factory A (Sensitivity Coefficient $\approx .70$) provides further reinforcement of the accuracy of the two models, even though the data used to build them came from very different sources.

More interesting, however, were the results for the Holder Cycle Time. It was the belief of the ingot plant engineers at the beginning of this research that the plant was constrained by the utilization of the Pits. This is expected since there are two Holders for each Pit. One would expect the single Pit to be continuously kept busy by the two furnaces and therefore be the bottleneck resource. However, near the end of this research, unrelated work by the plant engineers on another issue discovered that this belief may not be correct.

In fact, this research indicates that the Holder furnaces are the bottleneck resource as there is actually a significant amount of time spent waiting for the furnaces to be ready for the next Drop. This furnace waiting time exists because of the time spent filling the Holders with molten metal from the smelting operation and because the equipment operators focus on getting the Pit returned to production as quickly as possible. The model for Factory B confirms this new finding of the ingot plant engineers because the Sensitivity Coefficients are larger for the Holders than the Pit. This means that the biggest opportunity for improving throughput in the plant can actually be achieved by focusing on Holder Cycle Time rather than improving Pit operations. This conclusion can be used not only to modify behavior of the equipment operators but also to guide future capital investment.

4.2.2.2 Non-significant Inputs

Finally, examining the Cycle Time Variation numbers of both Factory A and B shows that in both models changes in Cycle Time Variation result in essentially no change in output. Variation is entered into the model as the coefficient of variation (standard deviation divided by the mean). The fact that this should have negligible effect on output is not immediately clear. One expects that high cycle time variation should lead to a decrease of output. Discussions with the software vendor have theorized that because of the aggregation inherent in the model, variation in cycle time is "averaged out" over long time periods. Another theory is that within the ranges over which the variation inputs were changed the increased cycle time variation has little effect.

4.3 Chapter Summary

This chapter presented the process for performing sensitivity analysis using the modeling tool described in Chapters 2 and 3. As this is software is a custom tool, it was necessary to describe the methodology used to obtain sensitivity values. The chapter concluded with a number of the important findings discovered as a result of the sensitivity analysis. This analysis served not only to improve the quality of the ingot capacity models but also as a basis for the capacity forecasting scenarios described in Chapter 5.

Cycle Times:													
Melter Cycle Time	20.0%	-14.0%	**-0.70**	-15.5%	**-0.78**	-7.8%	**-0.39**	2.5%	0.12	5.9%	**0.30**	5.9%	**0.29**
Holder Cycle Time	20.0%	-1.2%	-0.06	-1.7%	-0.08	-6.8%	**-0.34**	-9.0%	**-0.45**	-5.3%	**-0.27**	-4.6%	**-0.23**
Pit Cycle Time	20.0%	-1.1%	-0.05	-0.4%	-0.02	-6.4%	**-0.32**	-4.7%	**-0.24**	-7.7%	**-0.39**	-7.2%	**-0.36**
Cycle Time Variation:													
Melter Variation	50.0%	-0.1%	0.00	-0.1%	0.00	-0.2%	0.00	-0.4%	-0.01	-0.4%	-0.01	-0.3%	-0.01
Holder Variation	50.0%	0.0%	0.00	0.0%	0.00	-0.1%	0.00	-3.5%	-0.07	-1.8%	-0.04	-1.7%	-0.03
Pit Variation	50.0%	-0.1%	0.00	0.0%	0.00	0.0%	0.00	-0.1%	0.00	-0.2%	0.00	-0.1%	0.00
Drop Size:	-20.0%	-6.9%	**0.35**	-7.3%	**0.37**	-15.8%	**0.79**	-16.7%	**0.83**	-14.9%	**0.74**	-13.9%	**0.70**
Campaign Size:	-20.0%	0.1%	-0.01	0.0%	0.00	0.1%	0.00	0.0%	0.00	-0.3%	0.01	0.6%	-0.03
Melter Downtime:													
MTTR	50.0%	-0.6%	-0.01	-1.4%	-0.03	-1.1%	-0.02	-0.6%	-0.01	-0.5%	-0.01	-0.5%	-0.01
MTTF	-50.0%	-1.5%	0.03	-2.6%	0.05	-1.8%	0.04	-0.9%	0.02	-0.7%	0.01	-1.1%	0.02
Sched. Down	20.0%	-2.7%	-0.14	-3.1%	-0.15	-2.0%	-0.10	-3.2%	-0.16	-2.7%	-0.13	-2.2%	-0.11
Holder Downtime:													
MTTR	50.0%	-0.8%	-0.02	-0.1%	0.00	-1.7%	-0.03	-0.9%	-0.02	-1.4%	-0.03	-1.5%	-0.03
MTTF	-50.0%	-0.8%	0.02	-0.1%	0.00	-3.0%	0.06	-1.6%	0.03	-2.3%	0.05	-2.6%	0.05
Alloy Changeovers:													
Conversion Cycle Time	100.0%	0.0%	0.00	0.0%	0.00	0.0%	0.00	0.0%	0.00	0.0%	0.00	0.0%	0.00
D/C Cycle Time	100.0%	0.0%	0.00	0.0%	0.00	0.0%	0.00	0.0%	0.00	0.0%	0.00	0.0%	0.00
D/C/W Cycle Time	100.0%	-10.3%	-0.10	-8.9%	-0.09	-10.5%	-0.11	-7.7%	-0.08	-5.9%	-0.06	-5.9%	-0.06
Tool Change Time	50.0%	0.0%	0.00	0.0%	0.00	0.0%	0.00	-0.5%	-0.01	-1.1%	-0.02	0.0%	0.00
Recovery													
Melt Loss	100.0%	-1.9%	-0.02	-2.7%	-0.03	-3.3%	-0.03	-3.7%	-0.04	-3.5%	-0.04	-3.5%	-0.03
Scrap	20.0%	-2.0%	-0.10	-2.0%	-0.10	-3.7%	-0.18	-2.3%	-0.11	-2.1%	-0.11	-4.0%	**-0.20**
Miscellaneous													
No Double Pours		-0.8%		-0.6%		-2.0%		-8.7%		0.0%		0.3%	

Table 4.1 Sensitivity Coefficients for Factory A

Table 4.2 Sensitivity Coefficients for Factory B

Chapter 5: What-If Analysis

Once a capacity model was validated and understood through sensitivity analysis it could be applied to the examination of proposed capacity issues. This is the area in which the software tool presented in this thesis is expected to be most utilized in the future. The capacity model can be applied to understand situations which have never been experienced previously in the plant and thereby extend the intuition of the engineers into the range of production output that is possible. This type of exploration is referred to as What-If Analysis because it endeavors to answer the question, "What if the following production scenario were used in the plant?"

This chapter describes several categories of What-If scenarios for which the software tool presented in this thesis can be utilized. Because of confidentiality concerns, the details of the scenarios and their impact in improving ingot capacity cannot be presented in this thesis. The broad categories of issues that the ingot plant engineers are interested in examining are described to give the reader an understanding of the types of real world scenarios to which this tool can be applied.

5.1 What-If Capability

The software was built specifically with this type of analysis in mind. After running an initial simulation, the software tool has a What-If facility that allows the user to make changes to input values in the database model in order to determine their effect on output. As adjustments are made, the software records these changes, thus allowing the user to make many revisions simultaneously. Once this is completed, the simulation can be run again. The new output results are then displayed side-by-side with the original results. This works well for simple process changes where, for example, a cycle time is reduced for a particular operation, (e.g. melt time).

This new, shorter cycle time can then be entered into the model and the change in output can be observed directly compared to the original results.

Many of the situations that the ingot plant engineers are interested in exploring, however, are situations that involve more than a few simple changes. The relevant situations in which this tool is ideally applied involve dramatic changes in the assignment of alloys to Complexes. This can involve major changes in the make up of the model, which are discussed in more detail in Chapter 6. This type of examination can best be thought of as What-Do-We-Run-Where Analysis. This requires decisions as to which Complex should be used for each product and whether a product should be produced in more than one Complex. However, as discussed in Chapter 2, this software is not an optimizer and cannot automatically shift production from one Complex to another to balance product loading and increase output. Therefore this type of examination requires insight on the part of the user and iterations to explore multiple production possibilities. Two situations in which this type of work was necessary are discussed in the next section.

5.2 Scenario Analysis

During the course of the research, several scenarios were proposed by the ingot plant management to investigate for future production plans. The purpose was both to understand the predicted amount of output in each scenario as well as which might offer the most potential improvement. This process assumes that the plant continues to perform in its nominal state (unchanging yields, cycle times, etc.). The scenarios can be divided into two major categories. The first scenario type assumes that the physical plant and production processes are unchanged but the demand for many or all of the alloys changes significantly. The second type of scenario examines changes to or additions of plant equipment to determine the effect on output.

5.2.1 Product Mix

This scenario is best thought of in terms of traditional capacity forecasts. The ingot plants examine future demand forecasts at several levels of granularity (i.e. monthly, quarterly, etc). but by far the largest effort is put into the annual demand forecast process. Predicted demand on a per-product basis is compiled by the marketing organization based on discussions with existing customers and pooling the expected demand. This is provided to the ingot plant as a demand for each unique ingot product. This is a very detailed, time consuming process.

From a higher level, the demand for aluminum is often treated by grouping many alloys based on similar composition and similar types of applications. For example, if the demand for aerospace-grade aluminum is expected to rise, there are a number of alloys that fall into this category and therefore can be expected to have higher demand. Other alloys may either have reduced demand or provide a lower profit margin and the plant may simply choose not to produce these products in favor of the higher margin products. The demand forecasts, therefore, alert the entire factory to the fact that the relative mix of each product type the plant will produce will be changing in the coming year. Figured 5.1 indicates a typical way in which the demand forecast may be summarized for purposes of discussion within the plant.

- Alloy Group A ➜ 30% Increase
- Alloy Group B ➜ 20% Increase
- Alloy Group C ➜ 25% Decrease

Figure 5.1 Sample Change in Product Mix

This scenario is the type of application that the tool in this study was designed to consider. A change in product mix is entered into the model by adjusting both the number of pounds demanded for each product type and by modifying, where necessary, the percentage of each product that is to be produced in each Complex (see Chapter 6 discussion). The output indicates to the user the number of pounds of output that can be expected from the ingot plant when trying to meet this demand as well as the number of pounds per product type.

Again, however, the software does not optimize production across Complexes. A common result from this type of simulation is that one Complex may have far more demand than it can complete while another Complex may be under utilized. It is then the responsibility of the user to recognize these situations and make choices based on the capabilities of each Complex as well as the groupings of alloys that are assigned to each Complex. This provides the ability for the user to examine many different plant loading possibilities quickly in order to determine which might provide the best combination of overall output as well as production of high margin alloys.

5.2.2 Capital Improvements

The other type of scenario that was examined during this study was the opportunity to produce ingot types in additional Complexes by purchasing new equipment. This includes additional tooling for Complexes that do not presently produce a given cross-section. Also, certain alloys require additional equipment to process the metal that is not required by most other alloys. Therefore, this equipment may exist at only one Complex. Figured 5.2 gives an example of a type of equipment addition scenario that was examined.

- Add new tooling size to 3PIT
- Add new filter to 6PIT

Figure 5.2 Sample Complex Capital Improvements

The purpose of these acquisitions is not only to increase total capacity on particular alloys but also to provide additional flexibility to move alloys from one Complex to another as necessary. As such, from the perspective of the software, capital improvements simply become another case of modifying the product mix. Whereas a particular product or group of products may have previously run only on one Complex because of a unique piece of equipment, after additional equipment purchases those products may now run on two or more Complexes. It is therefore the responsibility of the user not only to reflect these changes in the model but also to make the optimization decisions concerning what percentage of each product should be produced in each Complex. This simulation exercise is then very much the same as examining a new product mix.

5.3 Chapter Summary

This chapter described a number of possible future scenarios for operating the ingot plants and how the software tool presented in this thesis has been used in their analysis. While the usefulness of this chapter to the reader is reduced by the need to keep actual scenarios and actual results confidential, the chapter attempts to provide an understanding of the What-If analysis capability desired by the ingot plant engineers.

Chapter 6: Conclusions and Future Research Topics

This chapter presents the conclusions and lessons learned during the course of this research. The Conclusions section revisits the goals presented in Chapter 1 and discusses the appropriateness of this software in terms of its maintainability. The Future Research Topics section describes areas in which the tool can be improved by presenting some of the recommendations that were made to the sponsoring company.

6.1 Conclusions

The original motivation for this work was presented in Chapter 1 of this thesis. Now we return to those goals to examine which goals were achieved, which were not and why. In addition, the research uncovered important barriers in the usability of the software tool that hindered its use. These barriers are described in the Database Maintenance section and recommendations to overcome these barriers are presented in the Future Research Topics section.

6.1.1 Goals Revisited

Recall the three goals of the research from Chapter 1: 1) Improve ingot capacity forecasting capability, 2) Examine future scenarios that provide opportunities for capacity improvement, and 3) Build the capability to examine two separate ingot plants as a single virtual factory model.

The first goal was achieve in this research. Chapter 3 describes the results of the efforts to build ingot capacity models for the two ingot plants in this study. Chapter 4 presents the sensitivity analysis work using these models to better understand the impacts of several aspects of manufacturing. The second goal was also achieved and was the subject of Chapter 5. Although the scenario results had to be excluded from this thesis for confidentiality reasons, the software tool was applied the analysis of future capacity scenarios during the research. The third goal, however, was not completed. Due to a limitation in the software tool discovered during the research, this portion could not be finished. The Future Research Topics section describes what this limitation was and what can be done in the future to resolve this issue.

6.1.2 Database Maintenance

A major conclusion from this work is that building capacity models using this software presently takes too long. The bottleneck is not, as was originally expected, the collection and characterization of production data. Building the model database in software and making modifications to the model actually consumed the majority of the research time. For example, the Sensitivity Analysis work discussed in Chapter 4 required approximately one full week of work utilizing two computers. This includes only the time necessary to make the necessary input changes to the model and rerun the tool. It does not include compilation, analysis, and interpretation of the results. The What-If analyses done for the company described in Chapter 5 also took one to two weeks.

Consuming large amounts of time to obtain capacity forecast results is a secondary issue for academic research. However, for a real world high volume manufacturing environment it is unacceptable to spend days modifying a model to examine a new production scenario. During the course of this research, several issues arose for which this tool could have been used if the model had been completed. In these cases, decisions and recommendations were typically needed within 24 hours. Intuition and "back-of-the-envelope" calculations are being used today in the ingot plants to address these issues.

The lesson for other software modeling efforts is that it is not sufficient to focus on execution time alone. While this is an important issue, the time necessary to prepare for model execution can be just as important. Even without further modifications to the software, the inevitable improvement in computing power will continue to reduce software execution times. However, additional computing power will not address the number of model modifications that may be necessary or the level of difficulty a user may experience with a poorly designed user interface.

Therefore, the modeling effort should focus as much on the required user interaction time as on the execution time. The more painful the user manipulation task is the less likely it is to be used. If the model requires a day in order to make each modification, it will probably not be used and thus fail to fulfill its objective. The following section describes recommendations that were made to the company and the software vendor at the end of this research to address this issue.

6.2 *Future Research Topics*

At this stage, the capacity modeling effort described in this thesis is still very much a work in progress. A primary owner has been identified at the company and industrial engineers at several plant sites are also using the software. However, there are several critical areas in which further development can be pursued to greatly enhance the usefulness of this modeling tool. These changes effect both the software tool and the construction of the database model. Two of these issues are described below.

6.2.1 Data Manipulation in the Model

During the course of this research it was apparent that there were a number of ways in which the user interfaces with the modeling software could be improved. The initial effort was intended to provide a connection between the software package and several of the computer databases available in the two factories in order to automatically populate the model as needed. This would provide a quick way to update the model when changes in the factory occurred. However, as experience with the tool increased it became clear that this was not the critical barrier to effective use of the software. Ingot plant operations change slowly and process characterization data does not need to be updated frequently. Yields, cycle times, and process steps are, in most cases, largely static for years at a time.

The model input that does change frequently, however, is product mix. Since examining changing product mixes was one of the primary reasons for creating this tool it is important that the process for entering this data into the model be straightforward. This will allow many potential mixes to be examined quickly (e.g. one to two hours). Using the present design of the software, however, this process proved to be extremely time consuming. Model modification times were typically on the order of one to two days.

Significant time investment was required for both creating product assignments to the Pits and defining process steps for every alloy-Pit combination. Each product assignment requires an individual entry in the model database for every assignment of a product to a Pit. For the number of products in use this is on the order of 1000 entries. In cases where an individual product is manufactured in multiple pits, the sum of the percentages assigned to each Pit must add up to 100%. Since all of this information must be entered manually it quickly becomes a tedious, error-prone task. In addition, there may be cases where user wants to investigate running alloys on Pits where they have not run previously. In these instances, since the processes are described separately for each pit, this requires building a tremendous amount of characterization data each time a new alloy-Pit combination is introduced. These two factors make the examination of new product mixes difficult.

Two ideas surfaced near the end of the research to make this job much simpler. The first idea was the realization that the complete set of alloy-Pit combinations that the user would like to examine is already entered in the user-specified Casting Plan (see Chapter 2). The task then became one of making the software intelligent enough to incorporate this information in order to build the required process characterization data for all alloy-Pit combinations. In order to do this, the process steps (see section 3.1.4.1) were reexamined to see where condensing of the data was possible. We determined that the cycle time required for each process step depends, to the first order, either on the alloy being run or on the Complex on which it is run. Two database tables were built, one with all the alloy-dependent process steps (e.g. cast time) and one with all the Complex-dependent process steps (e.g. melt time). We wrote a simple program to dynamically build all the necessary process data for each alloy-Pit combination as determined from the Casting Plan. Not only does this greatly reduce the amount of data that the model builder must maintain but it significantly reduces the chances for data entry error.

The second idea addressed the problem of creating the percentage assignments of products to Pits. Ingot plant engineers are seldom interested in adjusting one or two products from one pit to another. They are typically interested in moving entire alloys or groups of alloys across Pits. In the present software environment this requires changes to hundreds of individual database records. During this research a proposal was made to the software vendor to add a capability to define alloy groups in the software. Then, using these definitions, the software would need the additional capability to allow the user to set percentage assignments for each alloy group and have these assignments take effect for every product in that group.

Both of the two ideas described above were explored for feasibility with the vendor and agreed upon. Both will allow for much greater ease-of-use in building new models in the future as well as modifying existing models. It is critical that sufficient attention be paid to the time investment required on the part of the ingot plant industrial engineers. These engineers are the typical users of the tool and they have many other responsibilities in addition to capacity forecasting. If the software is too difficult to use or requires a significant time to relearn each time the individual returns to the tool it will fall into disuse. Both of these enhancements should address the long-term viability of the software in the production environment.

6.2.2 Virtual Ingot Plant

One of the stated goals of this research was to create a single model to include capacity forecasting for both Factory A and B. Proof of concept has been established as both factories can be satisfactorily modeled using the tool. In fact, during the research a single model was built and run using preliminary versions of the two factory models. However, in this process a barrier was identified that presently prevents a complete model from being built.

This barrier is the software's lack of a method for handling the difference in alloy transitions at the two factories. As described in Chapter 3, the procedures for preparing a furnace to cast a new alloy are significantly different between the two plants. With the present software design it is possible to define unique alloy transition processes for each Pit.

The difficulty arises from the alloy transition matrix. This is the list of all alloy pairs along with the procedure required to transition from one alloy to another. The software was only designed to contain one such matrix. So if, for instance, the transition from alloy A to alloy B in Factory A requires a Drain&Clean while in Factory B it requires a Filter Box Vacuum, it is not possible to specify this difference. This issue has been described to the software vendor and the problem has yet to be resolved.

Once this hurdle is overcome, the ability to simulate the two factories in a single, integrated model can impact combined real output. It can be a powerful vehicle for reinforcing a shared vision among the employees of the two plants. This increased sense of shared responsibility can drive behavior.

The area where the use of virtual factory model can provide the most significant benefit is in driving the development of shared practices. At present, many of the operations of the plant are significantly different. Much of this is driven from two major differences: 1) The two plants have different equipment with varying levels of automation, 2) The presence or absence of the Melter. These differences are understandable based on the way in which the two plants have evolved over the decades. These hardware differences justify many of the procedural differences that exist between the two plants.

However, procedural differences exist that could be addressed to drive synergy and improve operations at the two plants.. The alloy transition procedures provide a good example. These are operations that are not significantly impacted by the differences in production equipment. Much of the variance between the plants is driven by the presence of a Vacuum Crucible in Factory B to vacuum out residual metal from the furnace before a new alloy is added. The engineers in Factory B are committed to this practice and believe it is a crucial element in their efficiency and ability to quickly transition their furnaces. However, the engineers in Factory A whom have seen the procedure believe it is unnecessary and would not be helpful in their operations.

The potential benefits of building a virtual ingot plant go beyond the use of this modeling tool. The point of creating a "Virtual Factory" is two-fold: 1) Production can be seamlessly moved from one factory to another without loss in capacity or quality; 2) The factories in the network can learn from each other so that procedures in both factories improve more quickly than they would with each factory operating in isolation. During the course of this research many issues were observed where methods and practices vary widely between the two facilities.

The immediate virtual factory goal for this work is to create a single model for forecasting combined capacity for the two plants. However, the two plants ought to have a larger goal of achieving operational synergy so that manufacturing procedures are as similar as possible. The combined learning of the two ingot plants can address problems that have been solved in one facility but not in the other. This can be further formalized through the development of a shared process documentation system. This can speed the pace of the continuous improvement process occurring at both sites. Any help this tool can provide in beginning to build this shared mindset will benefit the overall operations of these two factories.

6.3 Chapter Summary

This chapter presented the overall conclusions from the research work performed for this thesis. Also, future issues were discussed describing areas where this research can be expanded going forward. This thesis has attempted to describe how mathematical modeling techniques can be combined with ever increasing computing power to analyze manufacturing in ways that were not previously possible. It is the hope of this author that this research will not only benefit the sponsoring company directly but also inspire future capacity modeling efforts in related areas.

REFERENCES

[1] Alcalde, Jeffery, *The Design and Implementation of a Synchronous Manufacturing System in a Job-Shop Environment*. LFM Thesis. Massachusetts Institute of Technology, Cambridge MA, 1997.

[2] Berry, Michael J. A. and Linoff, Gordon S., *Data Mining Techniques*. New York: John Wiley & Sons, Inc., 1997.

[3] Graham, Margaret B. W. and Pruitt, Bettye H., *R&D for Industry*. Cambridge: Cambridge University Press, 1990.

[4] Hatch, John E.., ed., *Aluminum Properties and Physical Metallurgy*. Metals Park, Ohio: American Society for Metals, 1984.

[5] Kann, Peter R. et. al. eds., *London Metal Exchange Prices*, The Wall Street Journal. New York: Dow Jones & Company, May 12, 1998, p. C18.

[6] Nelson, Paul E., *Development and Implementation of a Volume-Based Scheduling System for Highly Repetitive Production*. LFM Thesis. Massachusetts Institute of Technology, Cambridge MA, 1995.

[7] Senge, Peter M., *The Fifth Discipline*. New York: Currency Doubleday, 1990.

[8] Suri, Rajan and Sanders, Jerry L., "Performance Evaluation of Production Networks," *Handbooks in OR & MS*. 1993, pp. 199-285.